University of Kent at Canterbury

PROCEEDINGS OF

NETWORKSHOP 5

19th - 21st September 1979

Computing Laboratory

The University
Canterbury
Kent CT2 7NF
Canterbury 66822

PROCEEDINGS OF

NETWORKSHOP 5

19th - 21st September 1979

CONTENTS

Papers Not Included

The following papers arrived too late for publication:

Progress on Network Projects - RCO - Bill Hay

Preliminary Report on Job Transfer Protocol - Mike Guy

Reports on Machine Range Networking Activities - ICL 1900 - John Salter

INTRODUCTION

The papers in this document form a record of the proceedings of the Networkshop which was held at the University of Kent from 18th. to 21st. September 1979. This was the fifth in the series of Inter-University Workshops, which are now organised by the newly constituted Joint Network Team, (which is effectively a reincarnation of the Network Unit), which has been set up by the Computer Board and Research Councils.

The Workshop was attended by representatives from most British Universities, the Science Research Council and other institutions. A welcome development is the growing involvement of the Polytechnics. A record number of delegates attended and this properly reflects the growing importance of Networking and the associated standards.

A few personal comments on the meeting are perhaps in order.

1. Many projects are in the implementation stage and we must be patient in waiting for their fruits.

2. The need to standardise should perhaps be reiterated again, or more succinctly, "An enhanced subset of a standard is almost certainly not a proper standard".

3. The manpower shortage, which may intensify, will make even more co-operation essential.

4. The birth of "birds of a feather" sessions is a welcome development.

Finally, I would like to place on record my gratitude to my colleague Ian Dallas for carrying out the organisation of the conference so effectively and unobstrusively.

The next Workshop in the series will take place at the University of Durham in April 1980.

E. B. Spratt

University of Kent

October 1979

REPORT FROM THE JOINT NETWORK TEAM

ROLAND ROSNER

Joint Network Team

REPORT FROM THE JOINT NETWORK TEAM

Many of the topics in which the Joint Network Team has been involved over the past few months are covered at greater length elsewhere in these proceedings. The emphasis of its work has been to stimulate activity leading towards the development of components for networks. These include building blocks for local and wide area communications subnets as well as hardware and software packages for attaching machine range/operating system combinations to the by now familiar network hierarchy.

As indicated at the last networkshop, the production of such components must be carried out either by projects within our community or, given the shortage of expertise, under contracts with systems houses. In both cases, it is essential that the progress of a project be monitored by a steering group comprising those with appropriate expertise in the community.

A few projects have reached or are close to completion and details of some of these appear as separate papers. Among them are the Dataskil feasibility study on networking facilities for 1900's, the functional specification for the FTP on DECsystem 10's prepared by CAP and the X25 implementation for DECsystem 10's carried out by York University.

To put the funding of development work on a rather more formal basis, the JNT has presented a costed programme to the funding bodies and the indications are that, although the allocation of sums for expenditure on individual items remains to be negotiated, the outlines are agreed in principle.

The official date for the opening of PSS is still March 1980 and several centres have expressed their intentions to establish connections from an early date. It is expected that early usage of PSS will be somewhat experimental, the goals being to check out protocol implementations, to investigate the problems of offering services over a public network and to obtain reactions from a sample of users. Discussions with the Post Office on the provision of regional switching facilities at reasonable cost are taking place as a matter of urgency.

On the standards front the most interesting development is perhaps the activity which has been generated within CCITT on the Transport Service

Yellow Book. There are also moves to have it adopted as an interim standard by BSI. Implementation of the FTP is being undertaken at several centres and there will be discussion during the workshop of some of the deficiencies and difficulties encountered. As predicted when it was published, the Blue Book will undergo some revision to overcome such problems.

Between now and December, the Joint Network Team will grow to a strength of 4 with the recruitment of Barrie Charles, Ken Heard and Bob Cooper. Barrie is from the Rutherford Laboratory where he has been involved in data acquisition systems for experiments in high energy physics. Ken has been active on network projects at AERE Harwell for a number of years and is a leading participant both nationally and internationally in the development of networking standards. Bob has been responsible for the management and control functions of the British Steel packet-switched network. Two places in the JNT remain to be filled.

Dr R A Rosner
Joint Network Team

19 October 1979

REPORT FROM THE DATA COMMUNICATION PROTOCOLS UNIT

KEITH BARTLETT

Data Communication Protocols Unit

Data Communication Protocols Unit
- Progress Report

by K.A.Bartlett

Form and Function

The DCPU officially came into being four days after the end of Networkshop 3 at Cambridge and thus a report to Networkshop 5 comes conveniently at the end of its first year.

The Unit has a staff of only two full-time personnel including Peter Linington on secondment from the Computer Laboratory at Cambridge. However, the important thing about the Unit is its existence rather than its complement as it is more concerned to act as catalyst than to become deeply involved in all the development work. Most of the protocol developments are carried out in Universities, Manufacturers, Post Office and Government Research Establishments but the Protocols Unit is the mechanism through which these are related, reported and often initiated.

The Unit is providing a useful focus for all aspects of communication protocols and the UK continues to maintain a very strong position in the development and use of these protocols. The Unit is also being used as a substitute for the National Committee on Computer Networks (NCCN) as there is no other single point of contact in government for matters concerned with the development and use of switched data communication networks. The Post Office fulfils many of the required functions but not all and the Unit is used as a clearing house on a number of matters concerned with data communication techniques, developments and policy issues.

The Post Office encouraged the formation of the Unit. Support for the two Study Groups which remain active from EPSS has been split with the Post Office (Marketing Division) hosting Study Group 3 which is concerned with matters directly dependent upon the communications network. The Unit has taken on all network-independent activities previously addressed by Study Group 2 - the High Level Protocol Group. This formal division of responsibility has not affected the strong and effective working relationship between the two groups.

Protocol Development

We have concentrated on progressing a small number of functional protocols which will provide basic facilities over switched data communications networks - the spur to the work being PSS which is due to enter service early in 1980. This priority programme consists of four subjects:- Transport Service, File Transfer, Interactive Terminal and Transaction Processing protocols. During the course of this, other protocols have been identified and some work has also been carried out on a Job Transfer protocol.

Transport Service

The first activity of the Unit was to focus attention on the Transport Service problem through a small workshop in December 1978 which decided that the proposal made by a working party of Post Office Study Group 3 was worth further development. Consequently, the working party, chaired by Peter Linington, completed a detailed Transport Service specification which was printed and widely circulated (800 copies) by the Unit. Comments have been received from a large number of manufacturers and users both in the UK and overseas. These are being incorporated into a final specification which is being submitted as a British Standard. This is unusual in that British Standards are normally created following international agreement but here we have a suitable solution which cannot wait for resolution of the more difficult ISO problems.

The UK advantage in having a specification suitable for national use also enables us to set the pace in international discussions. This advantage accrues directly from the work of the EPSS community and we would like to think that it has been maintained because of the existence of the Protocols Unit.

File Transfer Protocol (FTP)

Again, the UK is in an excellent position due to work stimulated by EPSS. The original FTP specification was circulated early in 1978 by the High Level Protocol Group. This specification has been implemented by a number of sites who have formed an 'FTP Implementors Group' supported by the Unit. The several implementations constitute an effective trial of the protocol. The group continues to monitor progress and collates comments for incorporation into a revised specification of this basic bulk transfer protocol which will be used as a springboard for other activities - Job Transfer Protocol, for example.

While the present FTP is a valuable interim protocol and makes a valuable contribution to the international discussions, it is not being submitted as a British Standard. The long-term standard is seen as something more comprehensive in the form of a file transfer, access and management protocol.

Interactive Terminal Protocols (ITP)

PSS will offer a set of three related CCITT recommendations in this area - X3, X28 and X29 - and the very fact that they are PTT supported will cause these protocols to be replicated many times in equipments supplied for use with that network. Unfortunately, they have a number of serious shortcomings, such as an unfriendly user interface and not being compatible with a Transport Service. A specialist working party of Post Office Study Group 3 is studying the problems of 'triple X' procedures and has made recommendations on ways in which these protocols could be used over a Transport Service.

Transaction Processing

Many network activities will comprise the passing of a small piece of self-contained data from one party to another - possibly preceded by a short interrogation or authorisation sequence. The cost of these 'transactions' must be kept low and a study of Tramnsaction Processing is required to see how it might best be performed.

As with other protocol problems, the subject has been tackled by the Unit initiating a small special-interest group, chaired, in this case by a member of the DCPU Steering Committee. The group has not yet reported and it is therefore not clear whether transaction processing requires an individual approach or whether it is adequately covered by existing techniques.

Job Transfer Protocol (JTP)

Networkshop 3 made a clear call for work on remote job protocols such as a machine- independent Job Transfer Protocol to allow jobs to be initiated at one point, executed elsewhere in the network and the results returned to a nominated point.

The Unit has formed a working party on JTP. Unlike other working parties, this one will use the services of a consultant, employed part-time by the Unit, for the purpose of preparing a draft Job Transfer Protocol based on the FTP. A specification is expected early in 1980.

The formation of this working party, the use of a consultant to provide continuity and the mundane but necessary business of providing meeting rooms and refreshments are excellent examples of the purpose of the Unit which is to identify and facilitate protocol developments.

Problems associated with HLP's

Considerable 'missionary' work has been necessary to bring the developments in high level protocols to the notice of the wider but less-informed community of eventual network users (and even suppliers) so that their requirements and views may be incorporated into the standards. Unlike the Universities, many intending users of networks are not aware of the problems involved and the need for standardisation.

The international standardisation position is very fluid at present and it is much in the UK's interests to promote its view and present its proposals internationally. In general, our position is very strong and the detailed specifications on Transport Service and File Transfer form input to both CCITT and ISO meetings.

Support and Certification of Standards

When full international standards are agreed, it will be necessary to have test facilities to ensure that products are created in accordance with the specifications and that the full interworking which such standards promise is, in fact possible. The requirement for testing interim or national standards is similar but more urgent.

The Unit has drafted proposals for a National Protocol Certification Scheme which could issue certificates for manufacturer's interpretations of the current Transport Service and File Transfer protocols. This initial scheme must capable of extension to cover full validation of all high level data communication standards. One serious problem is that the standards are currently specified in natural language which is imprecise and leaves much interpretation.

For the present therefore, full compatability must rely on a certification scheme which is expert and strong enough to impose its interpretation as an extension of the specification in areas of doubt or ambiguity. In the longer term, the development of formal validation and protocol definition techniques must improve the specifications and simplify the testing.

	Transport Service	FTP	ITP	JTP	Transaction Processing
Working Group	PO SG3	Implementors Group	PO SG3	DCPU	DCPU
Specification	Early '79	Late '79	X3, X28, X29	Early '80	
Implementation		Several			
Final Specification	Early '80	FTP - YES FHP - NO			
Certification	Proposed	Proposed	PTT ?		
International Use	Probable	Possible	Yes		

DCPU - STATE OF THE PROTOCOLS

X25 LEVEL 3 P.O. TECHNICAL GUIDE

MIKE SANDS

Post Office Telecommunications

The presentation consisted of extracts from the Post Office X25 Level 3 Technical Guide.

For further information about this guide, please contact:

Mr. P. R. Morrison,
Post Office Telecommunications,
Lutyens House,
1/6 Finsbury Circus,
London,
EC2M 7LY

TESTING X25 IMPLEMENTATIONS

JOHN HORTON

GEC Computers Limited

and

JOHN THOMAS

South West Universities Regional Computer Centre

Developing and Testing an X.25 Interface

J R Horton GEC Computers Limited (GECCL)

J S Thomas South West Universities Regional Computing Centre (SWURCC)

1. Introduction

This paper describes the development and commissioning of an X.25 interface from the points of view of a manufacturer and a customer. A number of lessons learnt during the exercise are discussed in the hope that this experience will be of benefit to other potential network users.

The authors are proud to acknowledge the efforts and perseverence of the following people:

GECCL	Doug Steedman and Colin Watson
Honeywell	Bob Crawford
SWURCC	Pete White

2. Manufacturer's objectives and developments

When the Communications Systems Development group was established at GEC Computers Limited in 1977, one of its first objectives was to develop an X.25 interface for the GEC 4000 Series. It was rapidly established that there would be a number of difficulties associated with this task, particularly because of the lack of precision in areas of the CCITT recommendation which were left for further study and also because the British Post Office plans for PSS were completely unknown at the inception of the project.

As a consequence of these problems, it was decided to implement X.25 in a general purpose manner so that variations in the interpretation of the recommendation could be taken into account with little or no change to software or hardware. The product is based on a combination of system software and a programmed communications controller (PCC). The software is responsible for the user application interface, the entirety of Level 3 and the elements of procedure at Level 2. Communications line control and framing at Level 2 are assigned to the PCC which is loaded with a suitable microcode to accomplish this task when the system is initialised (See Figure 1). On the basis that software is simpler to construct and validate than microcode, this assignment of tasks has the advantages of placing those functions most likely to change within software and provides for a relatively stable controller microcode.

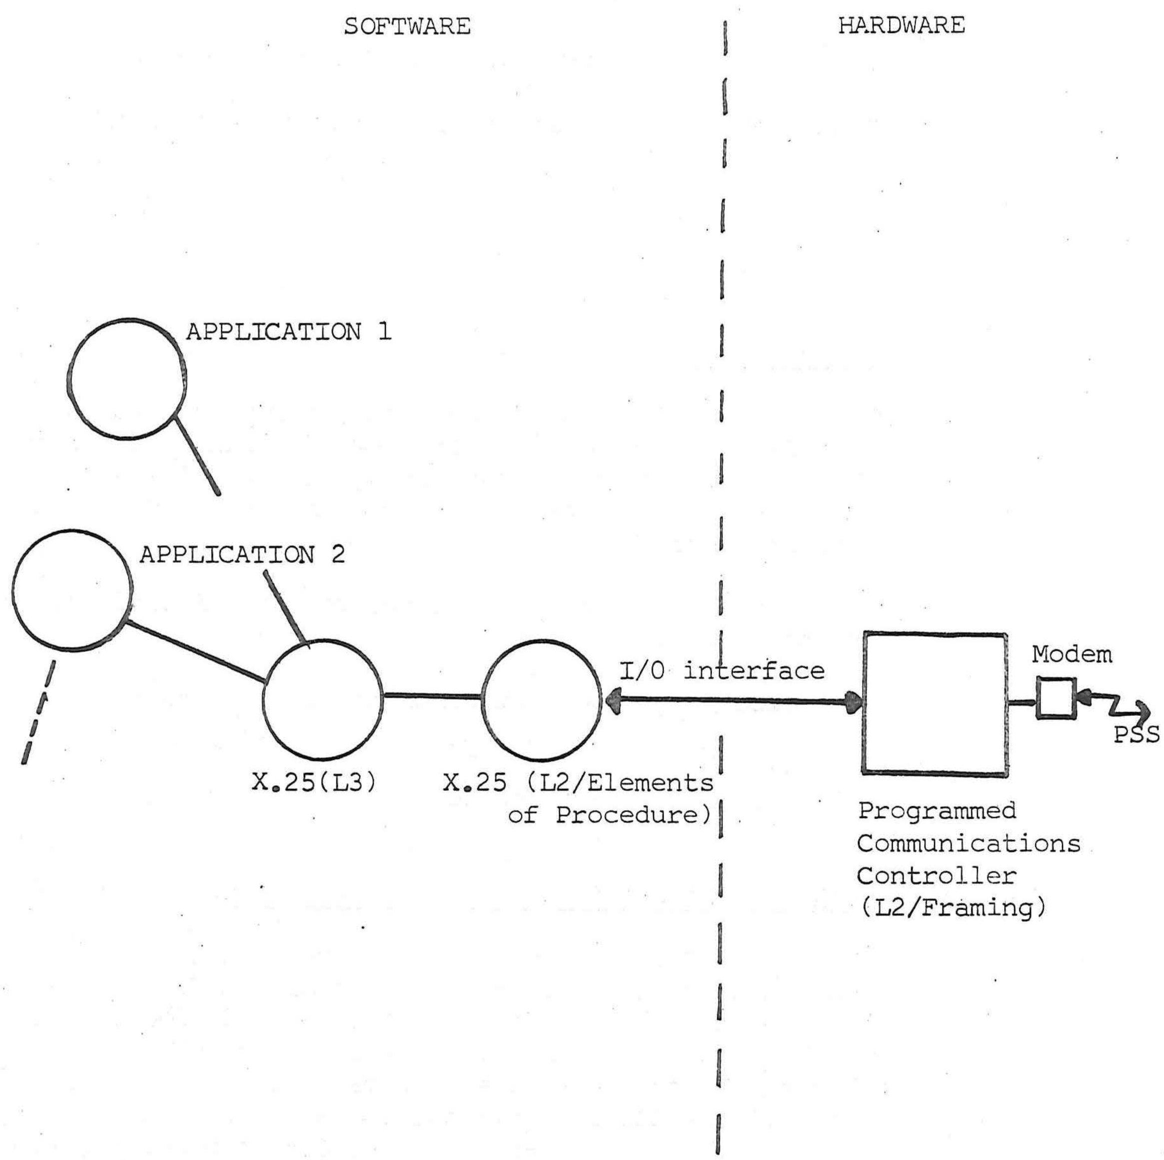

Figure 1 : GEC 4000 X.25 System Structure

As well as providing a full X.25 interface capable of operation at speeds up to 48Kbps, it was a further important objective to furnish the product with a good user interface. It was therefore decided to integrate the X.25 product with the Data Management system of the GEC multiple environment operating system OS4000. In this way a mapping between logical Data Management streams and X.25 Virtual Calls was established that allowed application processes to interface to X.25 in a natural way. A number of minor improvements and extensions were made to the Data Management system in order to accomodate functions such as call negotiation and the handling of interrupts and resets.

For the supply of X.25 in the UK, the Company took the view that the product would need to be compatible with PSS and would be available as both Data Terminal Equipment (DTE) and Data-Circuit Terminating Equipment (DCE). The Company was therefore prepared to take contractual commitments to meet the then unwritten specification for PSS with the proviso that all early releases of the product would eventually be upgraded to PSS compatibility to avoid long term support problems. This decision was taken in the light of two important factors. Firstly, it was a stated requirement of a number of procurement agencies that X.25 products should comply with Post Office standards. Secondly, despite their difficult position prior to the placement of the PSS contract, it became apparent that the Post Office was very willing to assist potential implementors to ensure that their plans and decisions broadly conformed with anticipated PSS developments. Thus when the draft Technical Guide for Level 2 was published, few changes needed to be made. At Level 3, although further work is required to implement Fast Select and Facilities, the software has already been structured to accomodate these features.

3. Customer requirements and constraints

An important component of the South West Network redevelopment plan is an X.25 link which will connect the Bath/Bristol replacement machine (a Honeywell Multics system) to the existing network via a GEC 4070. The latter will act as an interface processor between the new standard protocols being implemented on Multics and the existing non-standard protocols which the present network employs. The development plan required that Honeywell supply an X.25 DTE implementation to connect to an X.25 DCE package provided by GECCL. The protocol was required to meet PSS X.25 specifications and to operate on a 48Kbps link.

Figure 2 illustrates the software components present in the GEC 4070 Network Interface Processor (NIP) and shows the protocols being used above the X.25 package. In the particular case of the Multics connection, Honeywell chose to implement the X.25 package almost entirely in a front end processor although a Level 3 interface appears in the mainframe itself. The front end processor is a Honeywell Level 6 mini Computer and the final link configuration is shown in Figure 3.

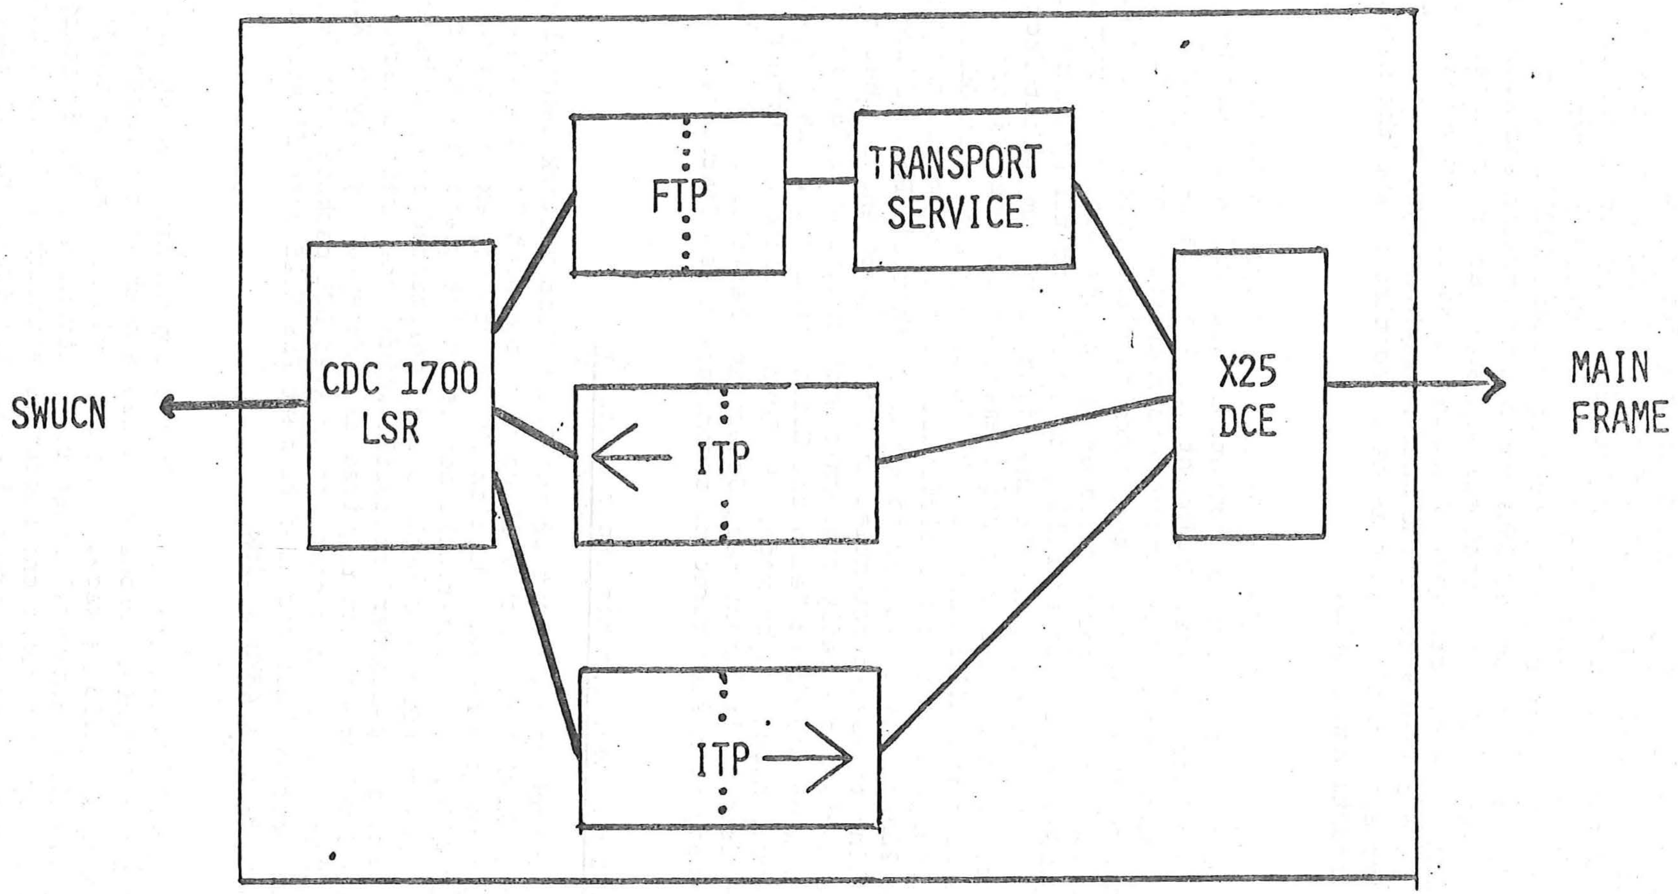

GEC 4085, 192 KB, CARTRIDGE TAPES

Figure 2 : Network Interface Processor

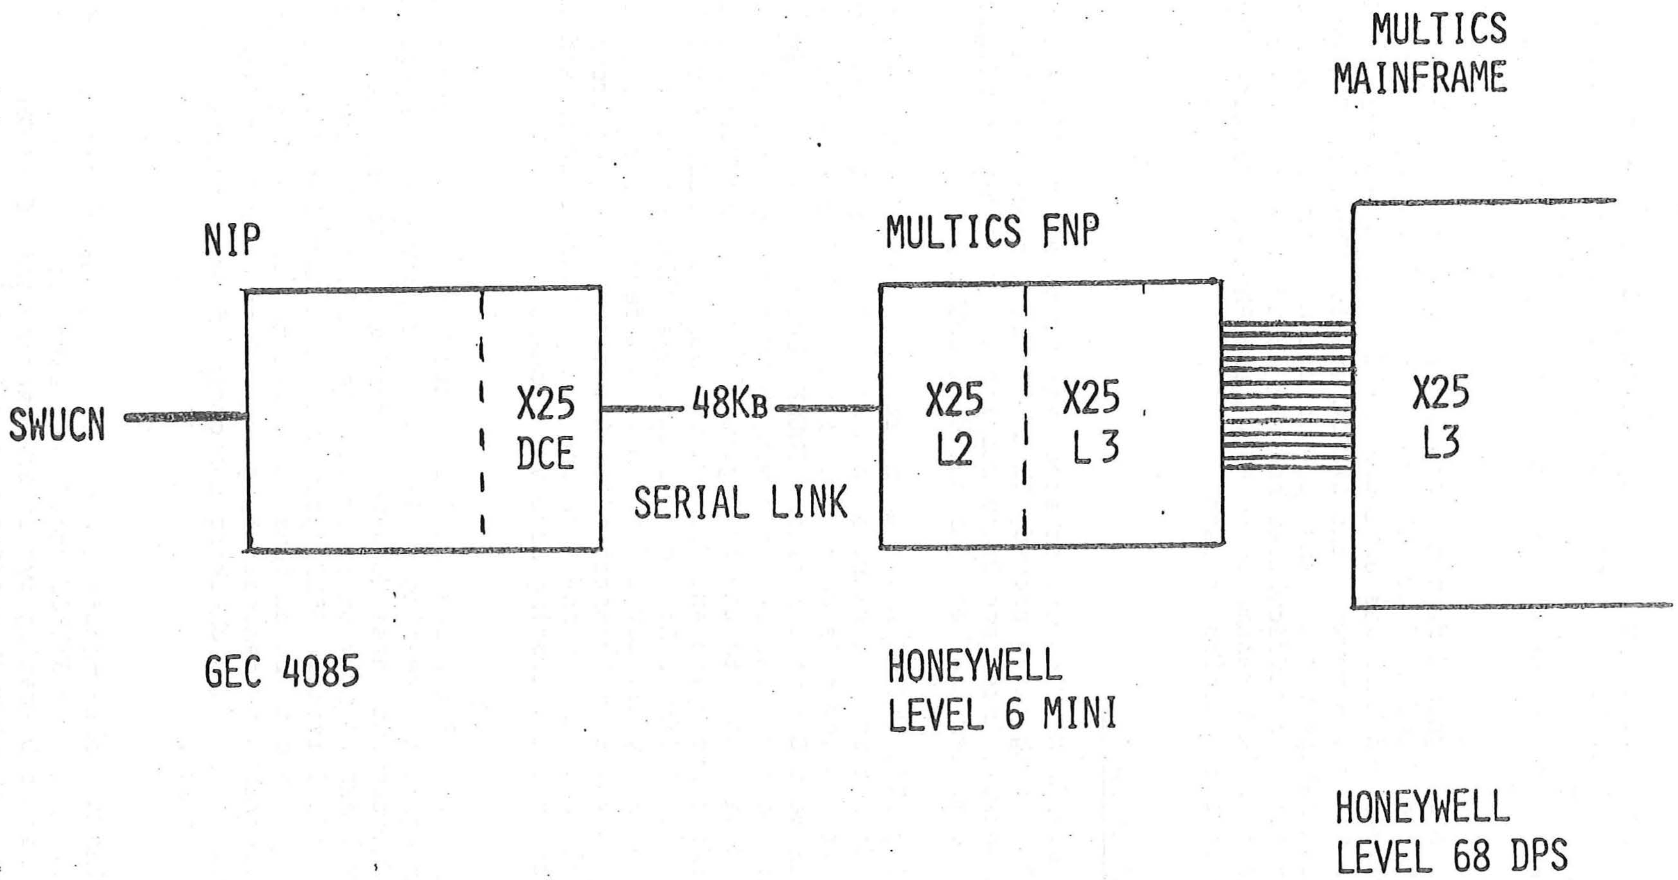

Figure 3 : X.25 Link Configuration

A separate development configuration GEC 4070 had previously been installed at Bath to enable the NIP software to be produced and a Level 6 Computer was also available locally when not fulfilling its primary rôle of providing RJE facilities for Bath users.

It was recognised from the start that the interconnection plan required above all else the wholehearted co-operation of all involved and this was achieved through joint meetings of GECCL, Post Office, Honeywell and SWURCC staff. The need for good personal relationships was all the more important bearing in mind the physical communication difficulties brought about by a complete system being developed in Borehamwood by GECCL and Phoenix, Arizona, by Honeywell for on site testing in Bath.

4. Project Review

The overall development strategy was produced in 1976 and by September 1977 funding was approved, GECCL had been chosen as the supplier of the Network Interface Processor and a development configuration GEC 4070 computer had been installed at SWURCC.

During 1978 SWURCC acted as a proving ground for GECCL's microprocessor based programmed communications controller, a prototype of which was used to implement a connection from the GEC 4070 to the existing South West Computer Network (SWUCN) CDC 1700 switch. By early 1979 this prototype was replaced with a production board and a further board was added to the configuration to provide X.25 framing, bit stuffing/stripping and Frame Check Sequence (FCS) handling at speeds of up to 9.6Kbps. At the same time a similar board was installed on the Honeywell Level 6 and both manufacturers supplied initial X.25 Level 2 implementations. After some teething troubles the basic problems of interconnection involving plugs, sockets and cables were sorted out and testing of Level 2 started.

By April 1979 the Level 2 implementations were thought to be satisfactory although later testing at 48Kbps was to prove this wrong. During the 9.6Kbps testing phase Honeywell staff were not present on site and SWURCC had to rely on their own diagnosis of the comprehensive Level 6 trace facilities. GECCL provided a Spectron Datascope line monitor which proved to be an invaluable aid to link level testing. This concentrated on demonstrating establishment and disconnection of the Level 2 link, and checking thoughput, flow control and recovery from link breaks.

During May both manufactureres supplied 48Kbps versions of their link controllers and SWURCC staff were also heavily involved installing the latest Operating System software release from GECCL and planning Level 3 acceptance tests and on site validation exercises.

From the outset the need to arrange for both manufacturers to be present on site for prolonged validation exercises lasting one or two weeks had been recognised and it was originally expected that one such exercise would be sufficient. Here again the assumption was proved wrong! The first such exercise took place during June when it was hoped to check both Levels 2 and 3 at 48Kbps. In the event only Level 2 was exercised and many bugs were uncovered that had hitherto remained undected; that they were not found earlier was a result of not having the two implementors together on site and also the increased speed of the link.

The opportunity was taken whilst Honeywell and GECCL were at Bath to evaluate their respective implementations for compatibility with the recently released PSS Level 2 Technical Guide. Both suppliers were subsequently to amend their CCITT X.25 implementations to conform to the new PSS X.25 standard.

At this time it was recognised that a further prolonged validation exercise was required and this was planned for the end of August. During July and early August effort was concentrated on the development by SWURCC of comprehensive X.25 Level 3 test facilities and the upgrading of Level 2 by the manufacturers to the PSS standard. Both manufacturers had by now supplied Level 3 software and basic on site testing was carried out by SWURCC as a build up to the second validation exercise. This exercise, which has only recently been completed, provided the opportunity to check out the now PSS compatible Level 2 fully and to build up the Level 3 testing with the aim of completing acceptance tests based on 60 virtual calls running on a 48Kbps link with pseudo-random resets and interrupts. X.25 software was also mounted on the production front end and NIP at Bristol so that the next phase of validation could begin. This will entail exercising the Multics main frame Level 3 interface as well as the front end and NIP software. The whole project is reviewed in Figure 4.

JANUARY	9.6KB HARDWARE
	LEVEL 2 TESTING
MAY	OS SOFTWARE UPGRADE
	TEST PROGRAM DEVELOPMENT
	INSTALLATION OF 48KB HARDWARE
JUNE	1ST ON-SITE EXERCISE
JULY	LEVEL 2 CHANGES (PSS)
	LEVEL 3 TESTING
AUGUST	2ND ON-SITE EXERCISE

Figure 4: X.25 Project Review (1979)

5. Lessons Learnt

During the course of the project a number of lessons have been learnt in connection with the development and testing of an X.25 implementation. The authors believe that these might be of assistance to other potential implementors and therefore present them below under a number of separate headings.

(a) Implementation of the X.25 protocol

At the framing level, the most serious problem was found to be ensuring that the order of data bits and their incorporation into the FCS was correct. The specification is obscure in this regard and a number of wrong implementations were experienced before a correct solution was obtained.

In the circumstances of project timescales, the implementation of a PSS compatible interface could only be achieved in stages:

(i) Test non-PSS compatible versions
(ii) Upgrade Level 2 to PSS compatibility after release of Level 2 Technical Guide
(iii) Upgrade Level 3 to PSS compatibility after release of Level 3 Technical Guide

Of these stages, only Level 3 Fast Select and Facilities remain to be upgraded. Future implementors should be able to proceed directly to PSS compatible systems.

(b) Pre-acceptance tests

It was found during initial testing that different test techniques cause different errors to manifest themselves. This is because of the critical dependency of the protocol and system software upon the timing and synchronisation of events. The following techniques are recommended:

* Internal software loop at Level 3 to validate test programs
* External hardware loop plugs to validate Level 2
* "Back to back" tests with two controllers on one machine
* Inter-machine tests to eliminate synchronisation effects

(c) Acceptance tests

For acceptance purposes, it is vital to have an independent test of the product and this is required at two levels. Firstly, for protocol compliance Post Office PSS test facilities will be taken to be the final arbiter. At this stage independent interpretation of the specification by two manufacturers and the customer was deemed to be a fair test.

Secondly, with regard to facilities and performance the customer or procurement agency would be well advised to provide a complex implementation evaluation test. In this case, such a test was provided by SWURCC and has proved invaluable. It is important to emphasise the difference between an implementation that can handle a dozen calls with a restricited protocol and limited user interface and one that is fully PSS compatible

and can handle upwards of 100 calls. The customer should set his acceptance criteria according to his anticipated future needs. A list of potential test criteria is given in the conclusion.

(d) Test procedure

During the project a number of surprises and pitfalls emerged:

* The faults found operating at 9.6Kbps were completely unrelated to those found at 48Kbps.

* A datascope class device is essential for debugging the link level and should preferably have both SYNC and HDLC/SYNC modes.

* Software trace facilities are essential but should be designed to be non-interfering. Small timing perturbations can mask problems. Large timing changes will affect the protocol operation.

* The achievement of a hardware connection is a non-trival exercise taking into consideration clock sources, modem eliminators, multiplexors, cross-over cables, plugs, sockets and V28/V35 electrical level conversion.

* The ability to test level by level is very useful although many inter-level problems will still be found. Software should therefore preferably provide access to Level 3, Level 2 and direct framing with potential loop-backs at each layer.

* Emulation of the end user, especially in a front-end configuration should not be neglected. Whilst data generation and looping is relatively simple, the provision of complete call control information is more difficult. The tester may be involved in a considerable amount of data entry at a terminal to supply this information. It may also be necessary to suppress protocol policing functions temporarily.

6. Conclusion

The main conclusion that can be drawn from this experience is that the provision of a broadly designed full X.25 implementation with a sound user interface, combined with a well managed integration and test project will provide the benefit of an operational X.25 system with few major difficulties. Attention to the completeness of the implementation reduces the number of potential incompatibilities and simplifies the incorporation of changes when specifications are announced. The integration project provides an incentive to get things working and gives a high degree of customer confidence in the final product when complete.

The cost of this approach is naturally high. From the manufacturer's point of view the implementation has taken between two and three man years with about four man months on site to perform tests. The customer, in addition to providing the test environment, needed to expend about four man months effort building systems, writing test software and controlling the validation exercise.

The authors suggest that considerable attention be given to the following criteria for the acceptance of X.25 systems on a cost/performance basis:

* Compatibility with PSS
* Completeness of X.25 implementation
* Integration with host operating system
* User interface
* Performance, considering:

  Packet throughput and delays
  Number of simultaneous calls
  Line utilisation efficiency
  Store occupancy
  CPU usage

In conclusion, SWURCC feels that it now possesses a very well tested product and that the two manufacturers and future customers will benefit from the exercise as well.

7. Acknowledgement

The authors are grateful to the General Electric Company Limited for permission to publish this paper.

REVISION OF THE TRANSPORT SERVICE

PETER LININGTON

Data Communication Protocols Unit

Revision of the Transport Service

P.F. Linington,
Data Communication Protocols Unit

Study Group 3 of the PSS User Forum has prepared the specification of a Transport Service, colloquially known as the Yellow Book. This specification was published for comment in April 1979; some eight hundred copies have been circulated and more than twenty written comments received, both from individuals and from organizations. This paper summarizes the updates intended as a result of the comments. For details of the original proposals, the reader is referred to the Yellow Book, copies of which can be obtained from the author.

Leaving aside changes intended to clarify the description, the changes of technical content fall into three areas: addition of two primitives to the service description, additions to the parameterization of the existing primitives, and clarification of the ways in which the Transport Service can be configured in different circumstances.

New Primitives

Address

One of the functions of the transport service is to manage the interpretation and transformation of the addresses used in establishing connections. Some applications, such as mail or task transfer, need to transmit addresses for later use in connection establishment. In order that the address manipulations are not duplicated and modularity is preserved, the transport service ADDRESS message provides the means of communicating addresses, transforming them as necessary.

Parameter form: ADDRESS address, qualifier

The message carries an address which may, at some future time, be used to establish a connection. If the address is valid in the sender's domain when the message is initiated, it will be transformed as it is transported so as to remain valid, and be delivered as an address valid in the receiver's domain.

The qualifier shows whether the message is currently passing towards or away from the object addressed. While the ADDRESS message is

travelling towards the object addressed, the transformations applicable to the called address in a CONNECT message are used. When the ADDRESS message is travelling away from the object addressed, those applicable to the calling address are used. The qualifier is changed if the message passes the nearest point on the connection to the addressed object. If an addressing error occurs, the fact will be indicated in the qualifier, and the message treated as if travelling away from the object addressed.

Interrupt

Because of the regulation of activity of the sending process and the attendent queueing implied by flow control, it may be necessary to modify or countermand previously transmitted requests still in transit. This requires an alternative means of communication allowing signals to be transmitted with priority, independent of the flow control constraints. Often, the ability to issue a RESET would suffice, but in some cases its action would be too destructive and other mechanisms, constructed using INTERRUPT messages, possibly in combination with DATA messages, must be used; these mechanisms will be application dependent.

The INTERRUPT message carries no parameters or data; it is an unqualified stimulus to which meaning is assigned by the application.

An INTERRUPT message is not subject to the same flow control constraints as data messages, and will normally be given priority over them. In all cases, an INTERRUPT message will be delivered before any subsequently transmitted DATA message. There will, however, be a separate flow control to limit the rate at which INTERRUPT messages can be transmitted so as to restrict the number that may be in transit.

The ways in which the INTERRUPT message will be used vary amongst applications, but the construction of common mechanisms, such as that to reset just one of the two paths, is for further study. Such mechanisms are outside the scope of the transport service.

Changes of Parameterization

The changes of parameterization proposed are confined to the CONNECT and ACCEPT primitives. In both cases, the change takes the form of the addition of a parameter indicating Quality of Service. In the case of the CONNECT message, the parameter indicates the Quality of Service

requested, while for the ACCEPT message, it reports the Quality of Service actually achieved in establishing the connection.

Configuration Strategy

A connection is characterized both by the primitive operations that can be performed on it, and by the quality of service it presents in performing these operations. The quality of service is concerned with the many intangible aspects reflecting the accuracy and efficiency with which the primitive operations are performed. For instance, the specification of error rates, bandwidth or cost effectiveness are matters of quality of service.

The quality required of a connection must be specified by the using application so as to fulfill its needs. Where the available quality falls short of the need, it must be enhanced. This can be achieved by optional enhancement functions provided in a network independent way; the network independence is ensured by both using and providing the same primitive operations at the upper and lower interfaces of the enhancement function.

Possible enhancements might be: cost reduction by means of multiplexing, reliability by means of error recovery or re-establishment of failing virtual circuits.

In cases where networks of markedly different quality are combined, the enhancement will often be on a network basis rather than an end-to-end basis. This may lead to the same mechanisms being used either within a network or across a network.

This strategy of using separate enhancements, both using and providing the Transport Service primitives, allows construction of a small number of network independent enhancement modules and the provision of only one network dependent Transport Service implementation per network.

Status

These changes are still being reviewed, and the Yellow Book proposals are being compared with those of other national and international bodies. It is hoped that a consolidated version will be published early in 1980.

HOW TO USE OLD KIT IN A MODERN NETWORKING ENVIRONMENT

TONY PEATFIELD

University of London Computer Centre

USE of OLD KIT in a NEW NETWORK ENVIRONMENT?

A. C. PEATFIELD - ULCC

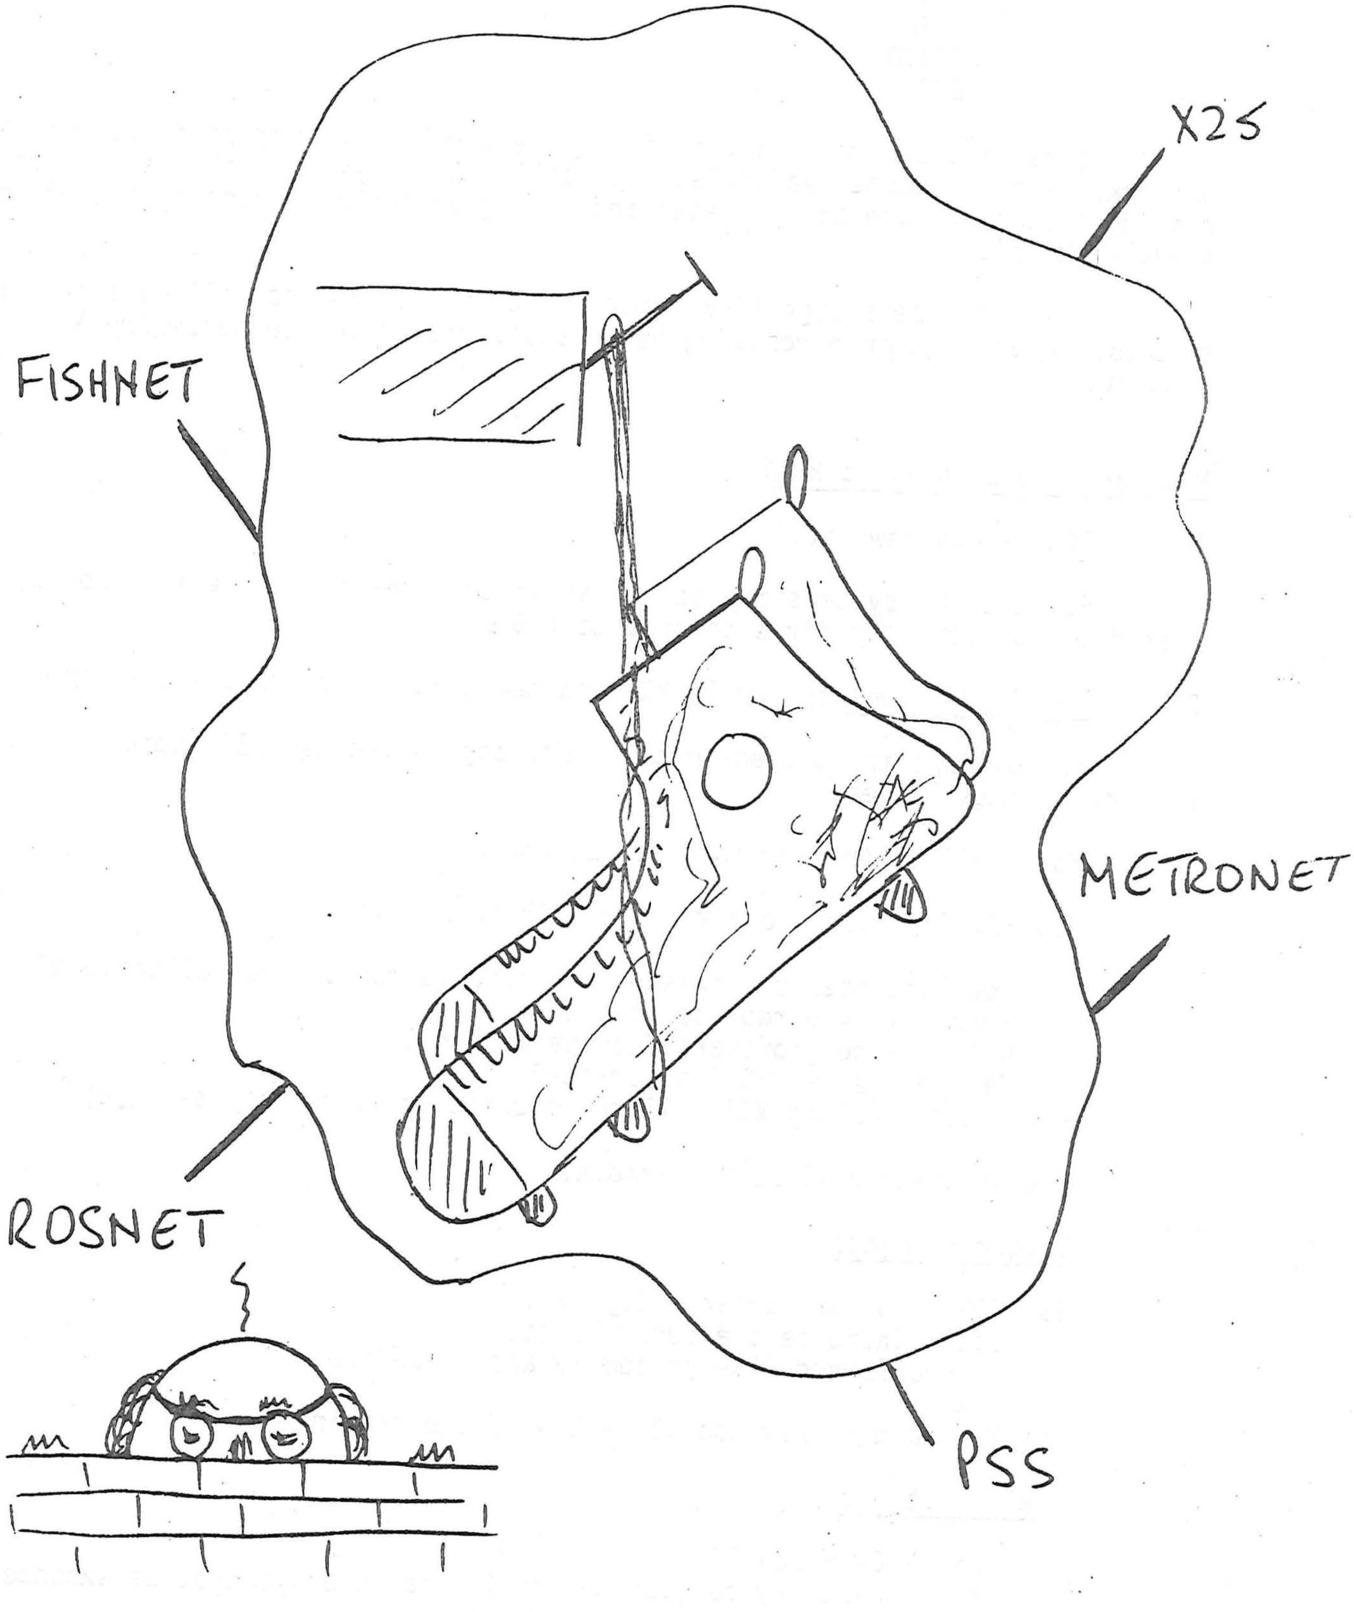

What is the old kit?

Users equipment presently comprises Remote Job Entry Stations and small computers with emulators of RJE stations.

They use manufacturer dependent protocols such as

2780
7020
HASP
UT200
ETC.

Users of the more comprehensive systems will want attaching to the new networks as they become available. Typically their systems will be big minis running a standard operating system and having an emulator running one of the above protocols.

There will be a significant portion of the users who will have to take this course rather than purchasing new systems that have the necessary X25 support.

How can we use this old kit?

1. DONT - Buy new kit.

But the new systems may be no better than the old. There is not a big rush from the computer firms to provide X25 support.

2. Put in New Code to handle X25 and new hardware to handle the line.

How big was the old emulator? - How big is the new X25 support? It may be 3 or 4 times larger.

For a PDP-11 the new code may be 16K or 24K.

Several Questions come to mind with this approach.

- How long does the present kit have to run before replacement?
- When will you replace it?
- Will the new software fit the machine?
- Can you get an HDLC interface?
- After handling X25 will the machine still do the users job?

Taking the PDP-11 as an Example.

Bigger Machines

Say PDP 11/34 or larger - Say RSX11-M.
- Could increase the memory size.
- You may spend £7K+ on memory and interface.

The X25 handler may use 10% - 20% of the machine??

Small Machines

Say PDP 11/04 - Say RT11.
- The machine may be just too small and memory cannot be expanded.

3. Put X25 Handler External to Machine

This could be a small system such as LSI-11 that can communicate with the user system by means of a simplified protocol. This system could of course also offer PAD facilities.

The cost of such a system could be similar to the costs of the portion of user system required to handle X25 if it is all done in the users system.

This solution may be the only one possible for smaller systems that are not capable of sufficient expansion.

Where should the X25 Function be performed?

In the case of host computer systems it is reasonably well accepted that the communications functions should be separated from the computing resources.

Should the same argument be applied at the user end of the network?

What do the Questions reveal?

There are a number of coordinated activities presently being pursued. These are aimed at producing X25 software for a range of computer systems.

There could well be a case for pursuing the other approach. Some systems will in fact need the external solution and a generalised system that could be centrally supported would be of great advantage.

This presentation was intended, hopefully, to do no more than make sure that everyone had thoroughly thought out all the available solutions.

TESTING HIGH LEVEL PROTOCOL IMPLEMENTATIONS

ALUN JONES

Computer Aided Design Centre

INTRODUCTION

This presentation is based on experience gained from the implementation of the Network Independent File Transfer Protocol at the CADCENTRE, and on a paper written for the FTP Implementors' Group (1).

BACKGROUND

The NIFTP implementation at the CADCENTRE was designed and written in the Autumn of 1978, and very quickly raised the questions:

a) How does one test an implementation of a brand new protocol?

b) How does one test a new implementation of an existing protocol?

The emphasis of these two questions is different in that the first is a means of testing the protocol, while the second requires the thorough testing of the implementation.

No guidelines were available in the NIFTP definition (2) as to how testing should proceed, apart from a few examples of the protocol mechanisms, given in section 7 of that document.

TESTING METHODS

There are two methods available to test a protocol implementation, namely,

a) Back-to-back testing, and
b) Site-to-site testing

The advantages and disadvantages of these methods may be summarised as follows.

BACK-TO-BACK

Advantages

1) Strictly controlled conditions
2) Any errors which occur must be local, the clearing of glaring bugs from the implementation may thus be done quickly.
3) Two processes may be tested simultaneously.

Disadvantages

1) Any misinterpretations of the protocol definition are likely be overlooked.
2) It is a narrow minded form of testing; the reason for implementing the protocol is to communicate with a remote site, and this should not be overlooked.

SITE-TO-SITE

Advantages

1) Avoids the problems of misinterpretation - any ambiguities in the protocol definition will come to light, and so be clarified.
2) Some error paths will be tested due to clashing implementations.
3) Produces a set of mutually compatible implementations.
4) Testing is done in the environment in which the implementation will have to work.

Disadvantages

1) Time could be wasted by one side having glaring bugs in the implementation, e.g. typing errors.
2) Only a single process is tested at any site at a given time.

From these summaries it can be seen that back-to-back testing is useful for the initial testing, while site-to-site testing should be used for the thorough testing of an implementation.

TESTING SCHEDULE

The testing schedule which was followed for the CADCENTRE's implementation of the NIFTP was divided into four stages, as follows:

Stage 1 - the back-to-back testing of level-ϕ of the FTP, the negotiation and termination level.

Stage 2 - the back-to-back testing of levels-1 and 2, the transfer control and data levels.

Stage 3 - initial site-to-site testing. This stage solved any problems of compatability.

Stage 4 - exhaustive site-to-site testing.

This schedule requires the implementor to provide both sides of negotiation and transfer from the outset. The thoroughness with which stage 4 is completed is dependent on the sites involved.

This schedule is reasonable when trying to test implementations of a new protocol where the number of sites is limited, but is likely to get cumbersome after a protocol has been accepted as standard.

VALIDATION

Ideally, new implementations of existing protocols would be tested against a validating implementation, which would follow a vigourous testing schedule to check the new implementation.

The objective of the validating site would be to provide a series of standard tests, and the onus would now be on the new implementation to prove its validity rather than to prove that it was not invalid, as is now the case.

Unfortunately, the provision of such facilities is unlikely to be an

easy task, and sites must be found which are willing and able to provide them. But they will prove to be necessary in the future if protocols are to be accepted as standard.

THE FUTURE

The testing of protocols and their implementations is not yet a thing of the past, and is likely to be a problem for the foreseeable future, since new protocols and applications are imminent.

At the moment, the study and practice of testing has not made any significant progress since the publishing of the NIFTP definition. It is still dependent on the goodwill existing in the small group of implementors.

There is a need for agreement on a formalised set of tests, and for the introduction of validating sites for the maintenance of standards.

REFERENCES

1) FTP Testing and Validation Schedule
by Alun Jones, CADCENTRE, Cambridge.

2) A Network Independant File Transfer Protocol,
prepared by the UK High Level Protocols Group, HLP/CP (78) 1.

PROGRESS ON NETWORK PROJECTS

- SOUTH WEST (JOHN THOMAS, S.W.U.R.C.C.)

- NORTH WEST (JOHN RICE, Liverpool)

- MIDLANDS (PHIL HARRISON, Nottingham)

- SRC (JOHN BURREN, Rutherford Laboratory)

SOUTH WEST UNIVERSITIES COMPUTER NETWORK

MID 1979 INTRODUCTION OF AUCC MULTICS

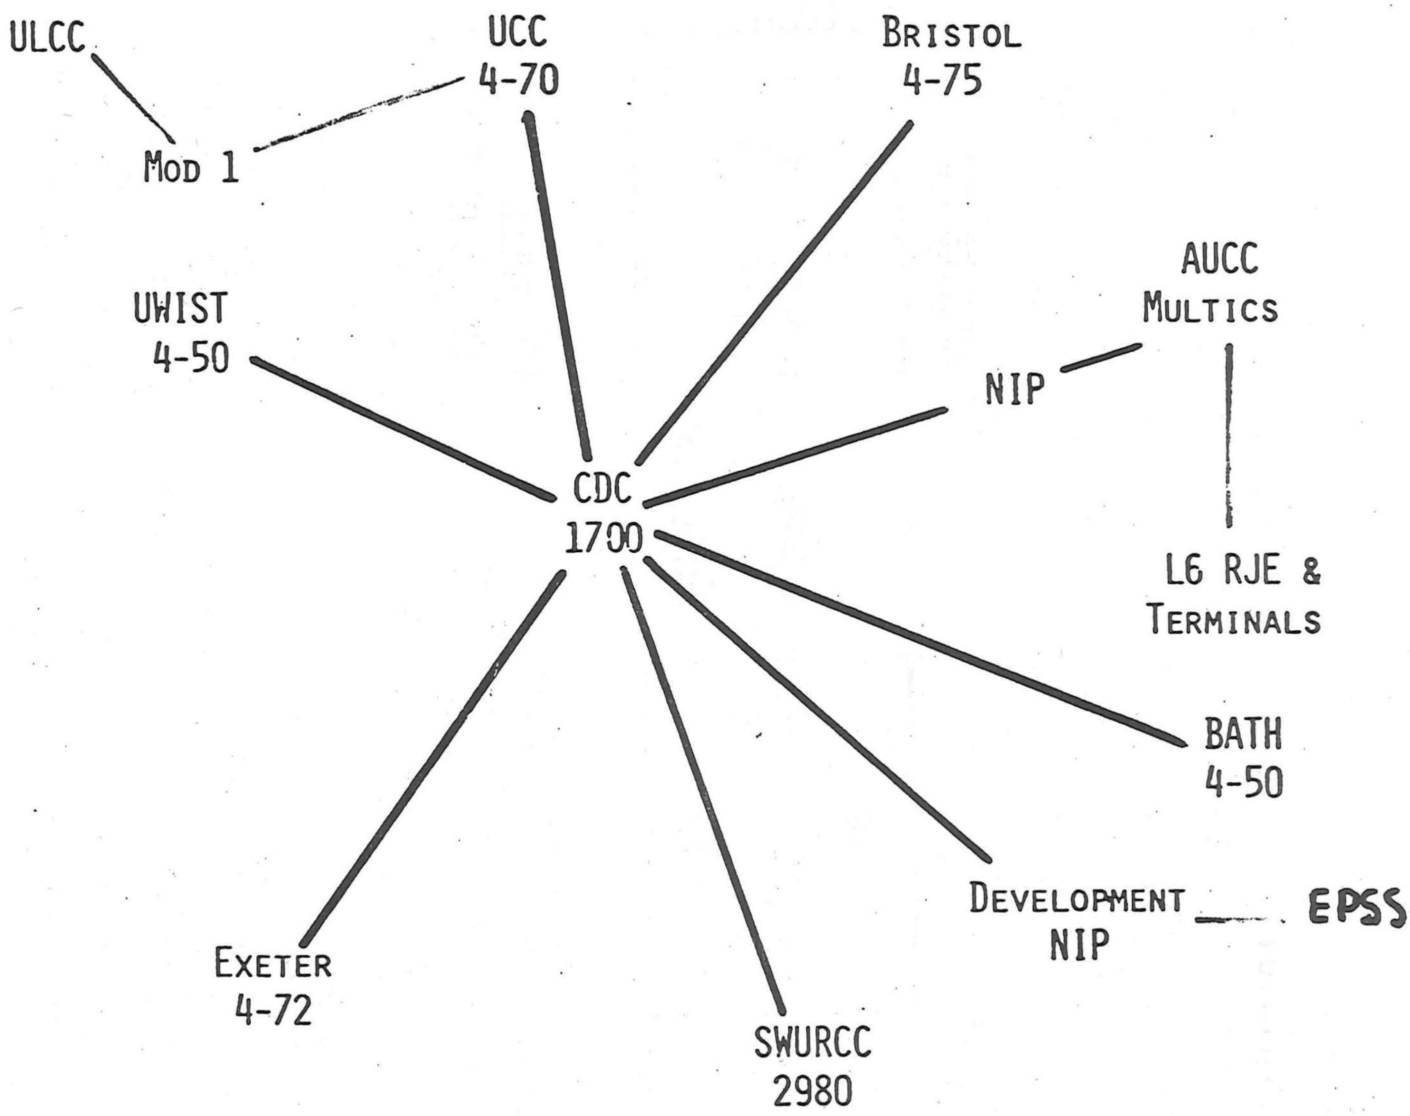

* AUCC - JOINT BATH/BRISTOL MACHINE
* INTERIM ACCESS CIRCUITS FOR EXETER, UCC AND UWIST TO AUCC
* NETWORK INTERFACE PROCESSOR
* 2980 BACK-UP LINKS
* EPSS CONNECTION

AUCC CONNECTION

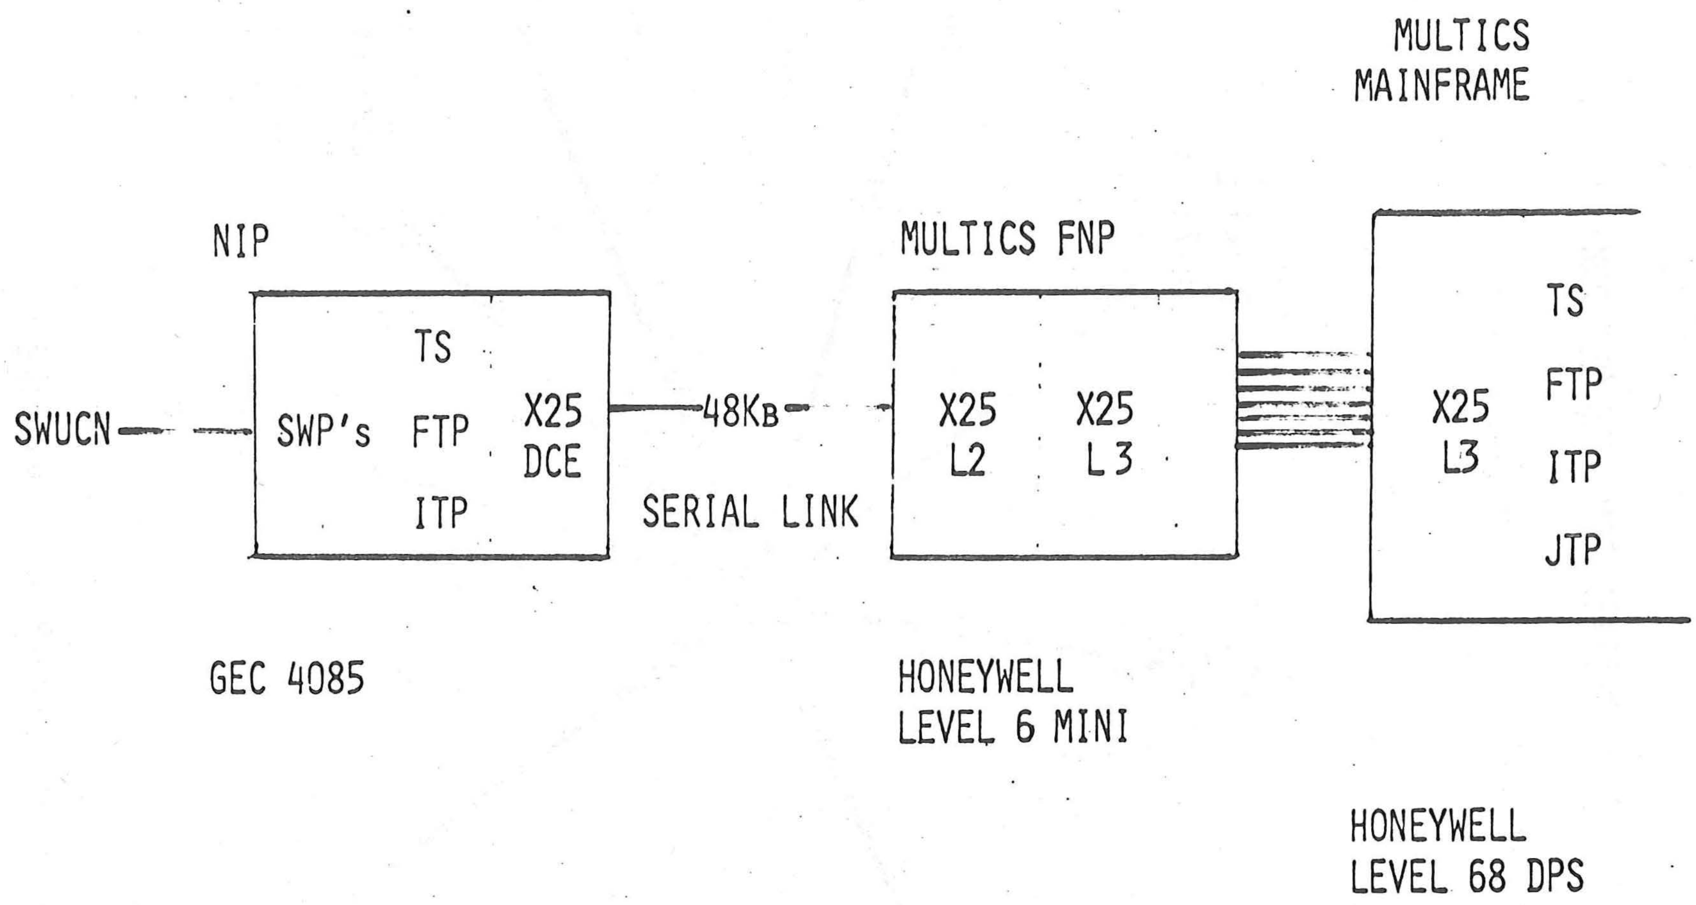

AUCC MULTICS CONNECTION

NIP		STATUS
SWUCN LINK, FTP & ITP	SWURCC	DONE
NI FTP & CONVERTER	"	DONE
X3/28/29 & CONVERTER	"	3/4 DONE
TRANSPORT SERVICE	"	DONE
X25 (DCE)	GECCL	DONE
LEVEL 6 FEP		
X25 (DTE, LEVEL 2 & 3')	HONEYWELL	DONE
FEP TO MULTICS INTERFACE	"	DONE
MULTICS MAINFRAME		
X25 (LEVEL 3' & FEP INTERFACE)	HONEYWELL	DONE
TRANSPORT SERVICE	"	START OCT 79
NI FTP	AUCC	3/4 DONE
X3/28/29	"	START JAN 80
JTP	"	STARTED

EUS via NETWORK LINK

SCHEDULED FOR JULY 1980

SOUTH WEST UNIVERSITIES COMPUTER NETWORK

MID 1980 SYSTEM 4S AT BATH AND BRISTOL REMOVED

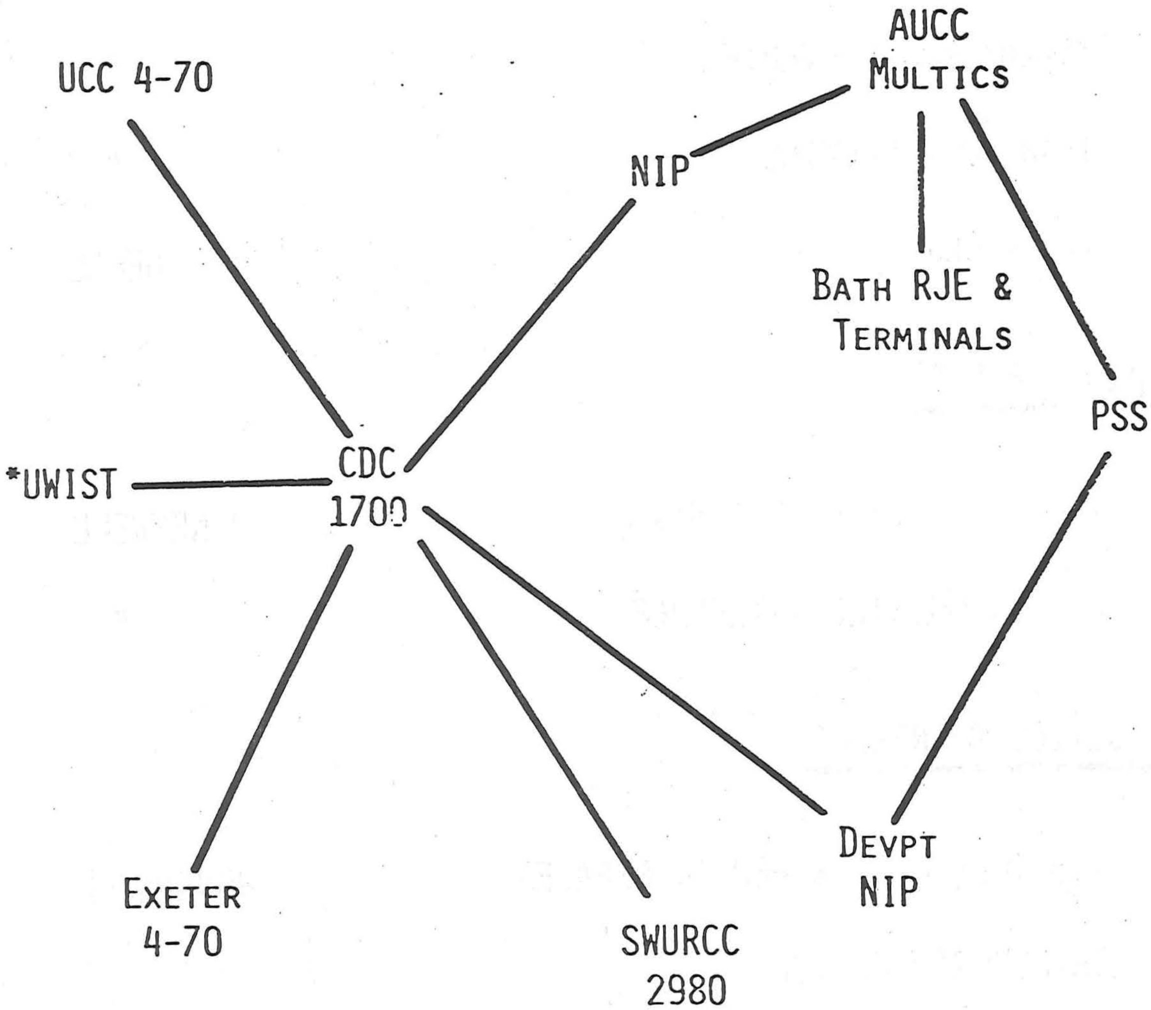

* SYSTEM 4S REMOVED FROM BATH AND BRISTOL (Summer 1980)
* FULL AUCC CONNECTION NOW OPERATIONAL " "
* INITIAL PSS CONNECTIONS MADE
* POSITION AT UWIST SUBJECT TO REPLACEMENT PROGRESS

SOUTH WEST UNIVERSITIES COMPUTER NETWORK

1981, 1982, 1983 1700 REPLACED BY REGIONAL SWITCH

Some mainframes will run links to both CDC 1700 & RPSE during transitional phase.

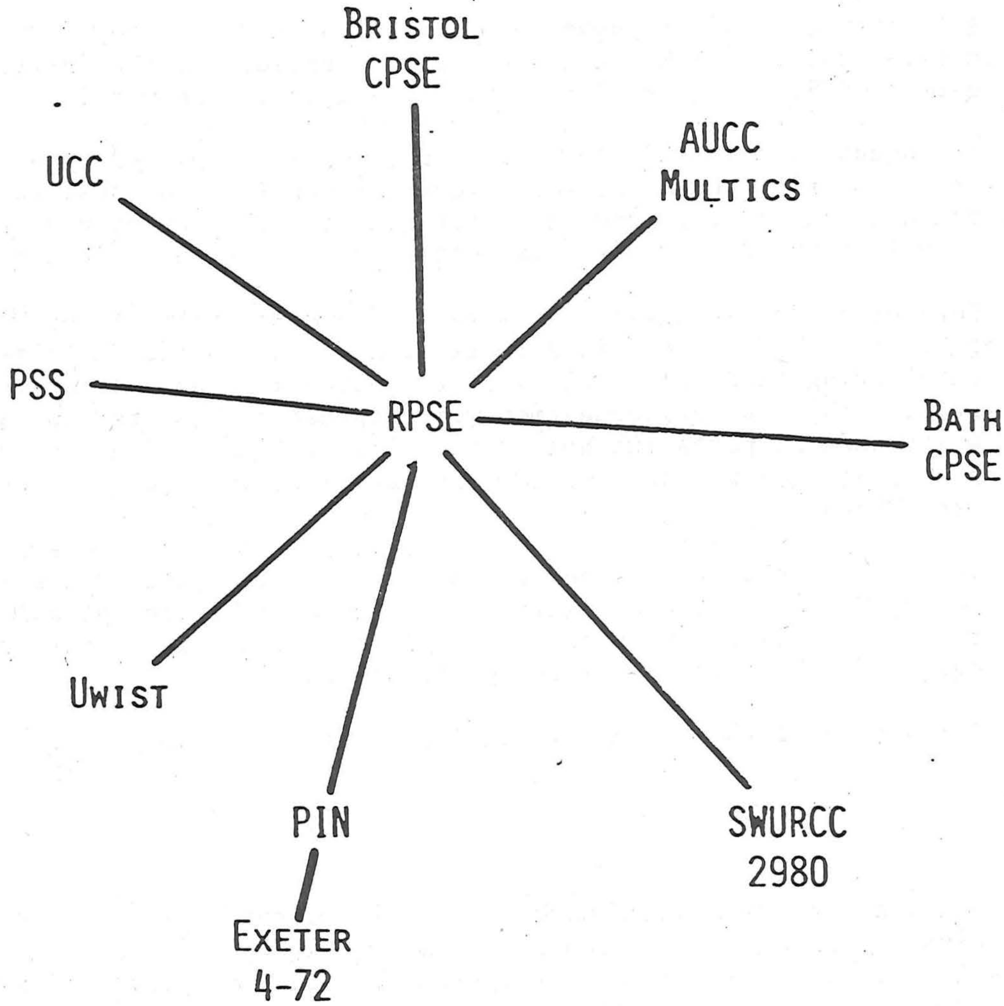

* CDC 1700
* RPSE/PSS SWITCH IN USE
* NEW MACHINES AT UCC AND UWIST
* X25 ETC AVAILABLE ON 2980
* REVERSE NIP AT EXETER
* CAMPUS NETWORKS AT ALL SITES

The North West Network :Current Status and Development Plan

J. D. Rice, University of Liverpool

Topology

The network topology at the commencement of user service in August 1977 consisted of nodes at Lancaster, Liverpool and Salford, at which were connected a 128K ICL 1905F, 512K ICL 1906S and 192K ICL 1904S respectively as hosts. All nodes support terminal connections, and an additional node was added at Keele in November 1977. A regular daily network service has been provided since January 1978 and since July 1978 a network service has been offered at each site from 10.30 a.m. until close of service.

A further node was brought into service at UMRCC in May 1979 to provide alternative routing for Keele users and as a prelude to the introduction of the hosts at UMRCC (ICL 1906A and CDC 7600) into the network.

In August 1979 the ICL 1905F at Lancaster was replaced by an ICL 2960 under DME. During the machine changeover a new issue of host software was introduced at other sites and Lancaster successfully re-entered the network with only 1 of the 30,000 words of host network code requiring amendment.

Current effort is directed towards the introduction of the UMRCC hosts during October 1979. Additionally it is hoped to connect a Prime 400 at Salford offering terminal services on the network in approximately the same timescale. This latter connection will represent the first instance of a site whose node supports two hosts, the first instance of a host connected to the node using a regional variant of HDLC LAP, and the first connection of a non-ICL host.

In January 1980 it is hoped to connect the Keele GEC 4082 as a network host, using X25 level 2 LAPB (which will then also be used at Salford) and in April 1980 it is expected that inter-node communication, currently using the regional LAP variant, will change to use LAPB.

Connection of the Lancaster ICL 2960 under VME/K is scheduled for late 1980.

Use

Detailed analysis of network traffic is currently in progress. However, preliminary figures are available for use of Liverpool from network terminals at Keele, Lancaster and Salford (figure 1). This displays the usual vacation-time troughs common to university computer usage. More interesting information can be obtained from the breakdown of the terminal traffic by source site (figure 2). Three distinct types of network usage may clearly be distinguished:

(i) General Service
- Lancaster's mainframe was removed at the beginning of June and its replacement commenced service approximately two months later. Then local hiatus corresponded with the highest remote site usage, which was preceded by a build-up as work was transferred to Liverpool in preparation.

(ii) Research Service
- Salford's network use is specifically oriented towards the research user whose particular computing needs are best met at Liverpool. This workload is reasonably stable.

(iii) Teaching Service
- Keele's network use is strongly related to teaching requirements as shown by the marked decline in use during vacations.

Further analysis of the network service use of Liverpool by Keele, Lancaster and Salford reveals that the average length of a terminal session is 17 minutes and that an average of 35 such sessions take place per day.

Facilities

The facilities offered by the network have now been completed on the network mainframes and provide at both the job control language level and application protocol level for:

filestore to filestore transfer
job transfer
and filestore to remote device transfer.

Considerable effort has been expended on the definition of a simple, acceptable user interface for the exploitation of these network facilities.

The key concepts behind this interface are those of:-
(a) remote relationship

and (b) output routing destination

The requirements were first defined at JCL level as:

(a) REMOTE HOST hostname,:username,PASSWORD password

(b) OUTPUT HOST hostname,:username,PASSWORD password, PROPERTY qualifier

and the implications of these definitions were then considered in relation to all existing commands and high level protocols.

The implications at the user interface level of the introduction of these new concepts were minimal, amounting mainly to the addition of one or two parameters to existing commands. However, a comprehensive set of defaults eliminated the need for change in the majority of operations (see Appendix 1 :NWD25).

Developments

Recent developments in the communication node software include

(a) an extension to handle multiple hosts at a node.

(b) an extension to the character terminal device protocol to allow "free running" terminal access (i.e. type-ahead).

These two developments were necessary for the addition of a second host at Salford, two hosts at Keele, and the ability of all of these new hosts to support type-ahead across the network.

Standards

The region has no commitment to the long term maintenance of existing non-standard protocols. As opportunity for revision has arisen, advantage has been taken to implement standard protocols.

(i) HDLC

the inter-node protocol is currently a variant of LAP.
The requirement arose to connect non-ICL hosts to nodes and the opportunity was taken to use LAPB for this purpose. It is necessary because of storage restraints to minimise the number of device

drivers used on each node, and therefore the inter-node protocol will also be changed to LAPB.

(ii) FTP

The introduction of non-ICL GEORGE 3 hosts presented an opportunity for revision of the existing file transfer protocol which sent files as a transparent binary stream.

In particular, a requirement for ASCII record transfer was identified. It was therefore decided to implement FTP-B for this purpose. A subset suitable for both the new requirement and the replacement of the existing requirement has been determined.

It is a requirement that any new protocol should enable the same service to be provided and the SRC -"subset" was found to define less than our requirement, principally in:

(i) the mode of access

(ii) restart markers

One new attribute has also been identified:

for the "qualifier" to "LP" in "take job output"

and a new error message indicating: "command valid but not executable immediately".

In October 1979 it is intended to submit this subset for official approval as a standard set of non-default options for FTP-B.

PSS

It is hoped that Liverpool will act as development site for a Board-funded project by ICL Dataskil to implement PSS connection up to the Transport Service within GEORGE 3 for ICL 1900's, with Liverpool validating built-in and built-out FTP against this Transport Service, by mid-1981.

It is hoped that Salford will be able to connect both their Prime and ICL 1904S to PSS and engage in the pilot implementation of JTP.

UMRCC intend to provide FTP and simple JTP facilities over PSS using a GEC 4070 front ending the ICL 1906A, by the end of 1980.

Conclusion

The region now has available a variety of services via the North West Network for its users. During the next four years it will be necessary to implement a smooth transition from private to community-wide protocols, and from private hardware to standard switches whilst improving the level and range of services.

Network Terminal Sessions at Liverpool :1979

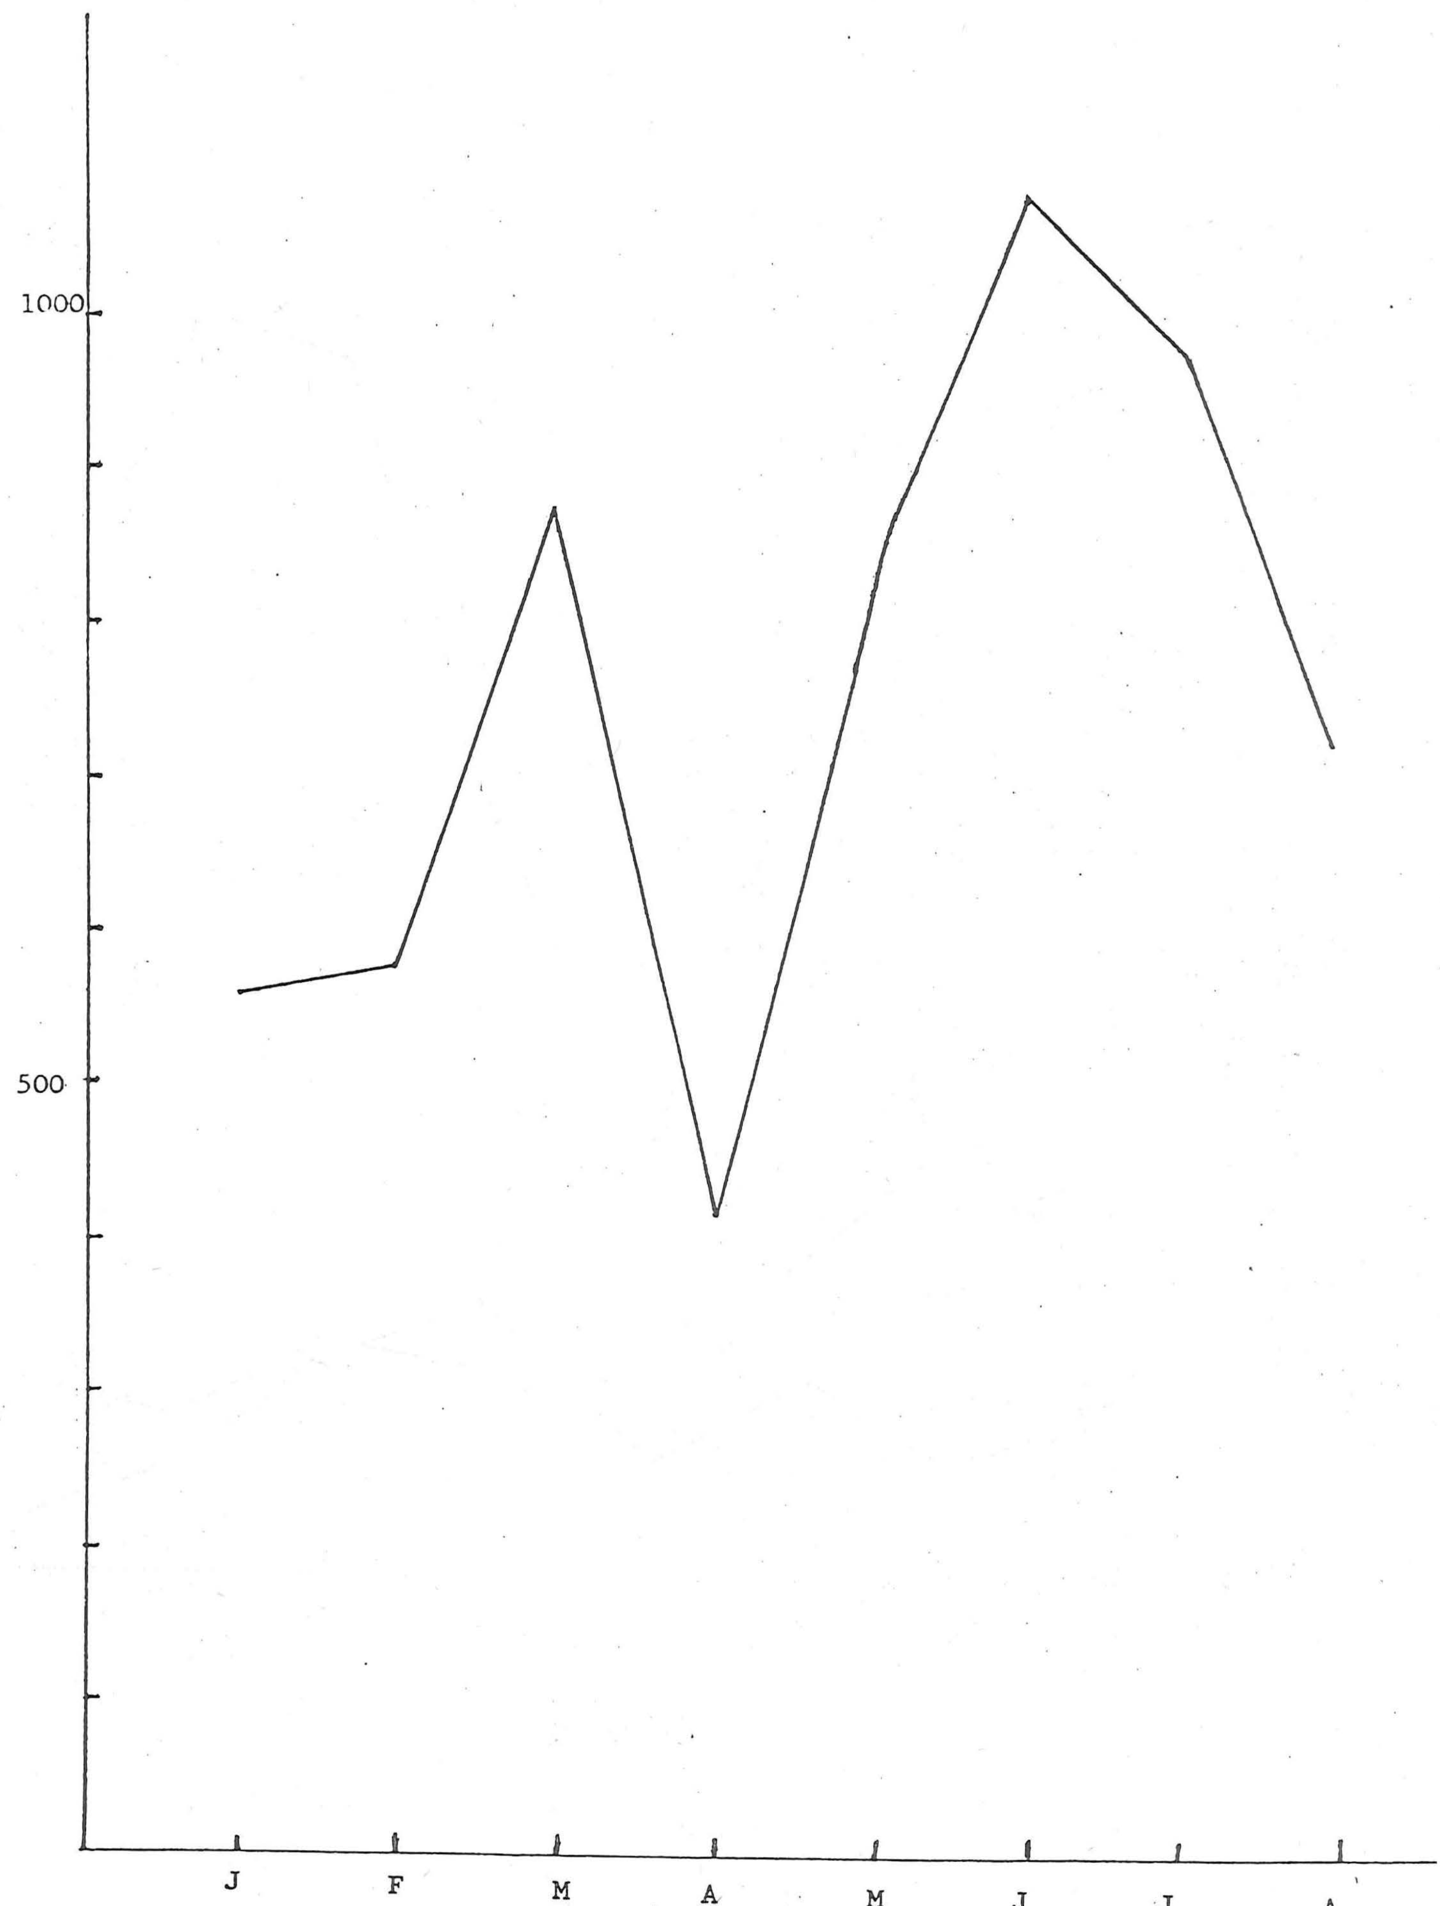

FIGURE 1

Network Terminal Sessions at Liverpool :1979

-by site of origin

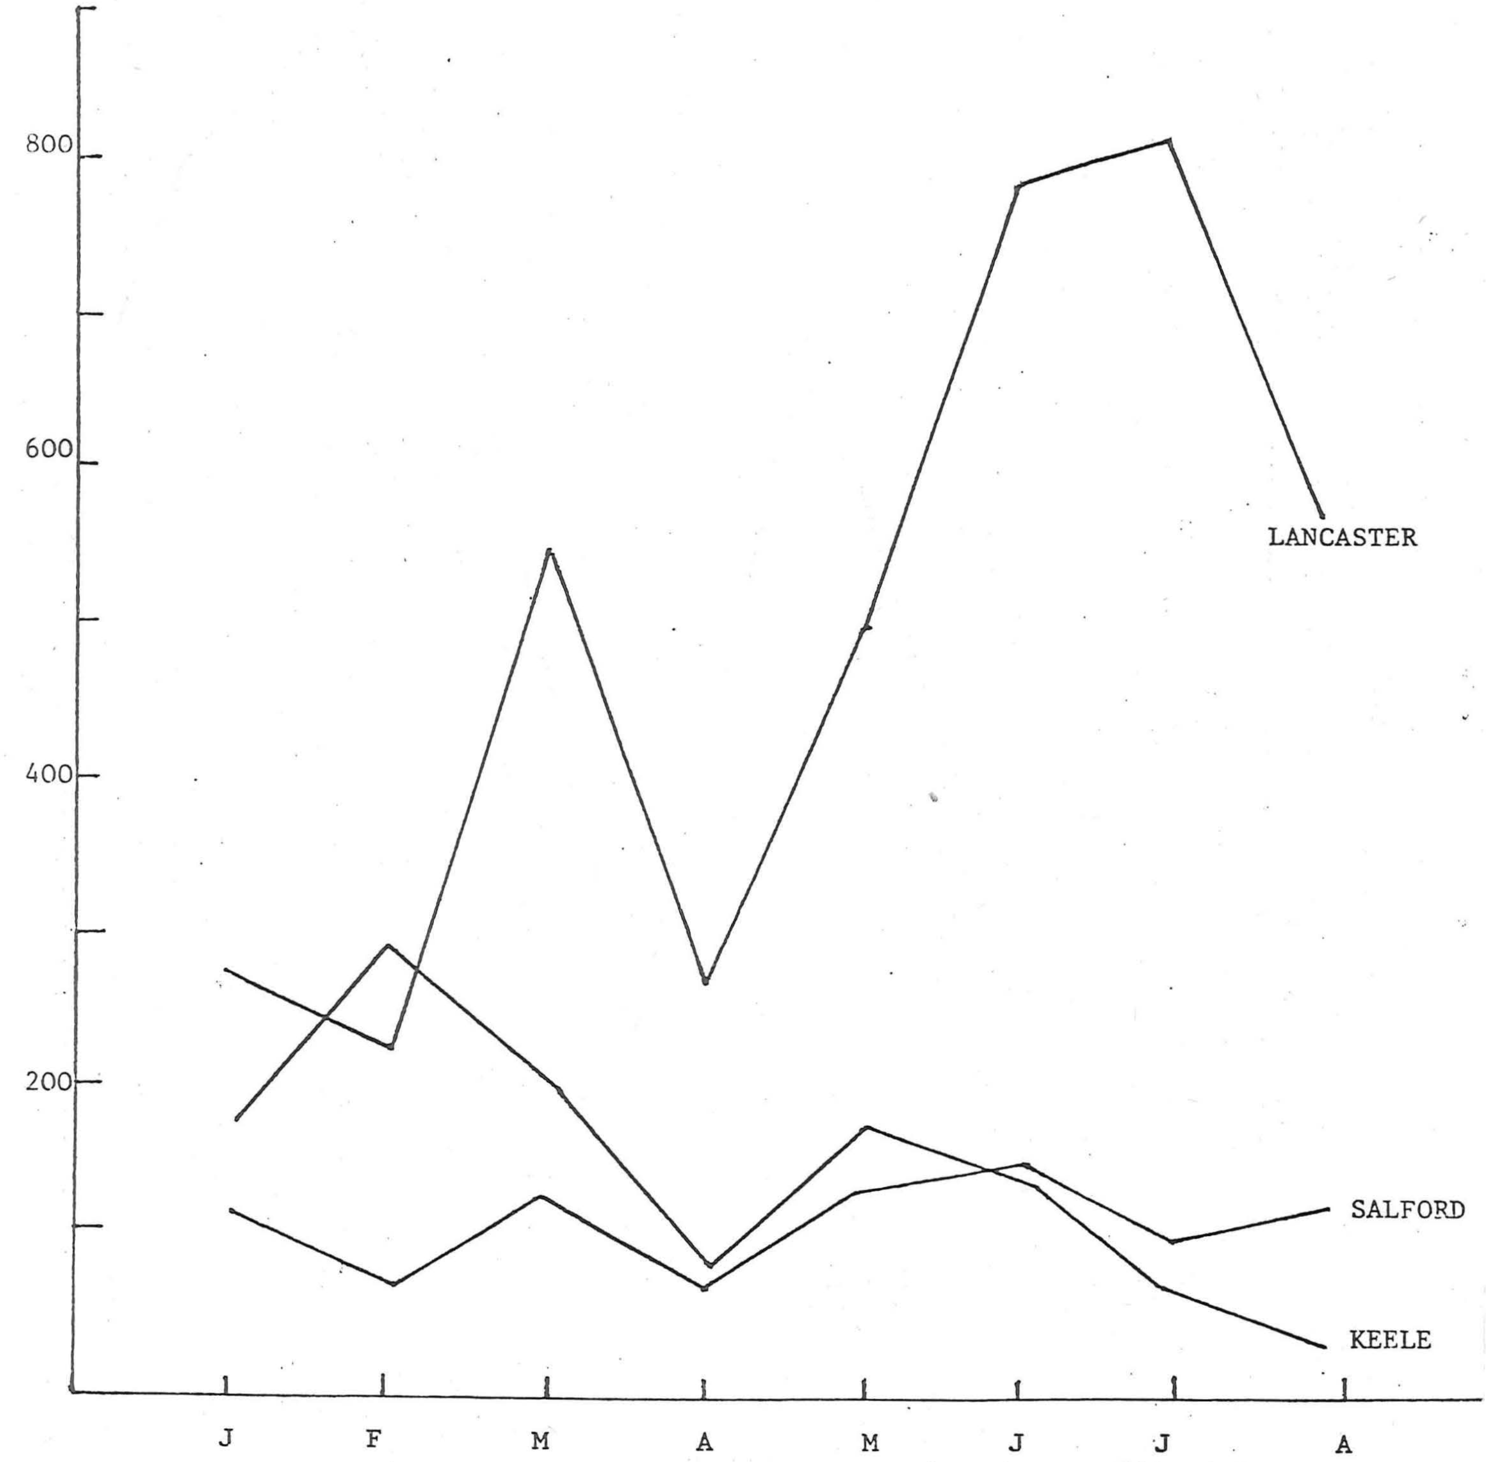

FIGURE 2

OBEY

OB *file*,PARAM(*param*,*param*,...)

Causes the macro in the file to be obeyed, with the parameters being used as required.

OUTPUT

OU HOST *host*,*username*,PW *password*,PR *property*

Defines a host site and associated property at that site, to which all output produced as a result of subsequent LISTFILE, JOB (in batch use) or RUNJOB commands is to be sent. *username* and *password* are those applicable at that site. If this command is not used, the host and property associated with the user's terminal or the input device used for a batch job will be implied, but this command must always be issued before any LISTFILE command that will cause a network transfer, unless a suitable OUTPUT parameter is included in the LISTFILE command or in the LOGIN command. The site and property can be changed for individual LF, JB or RJ commands by the inclusion of an OUTPUT parameter. OUTPUT without parameters obtains a report of the values currently in force.

PRINTLAST

PL

Causes the last record sent to the monitor file to be displayed. Can be used to obtain the full version of an error message suppressed by QUIET.

QUIET

QI

Causes command error messages to be replaced by ERROR.

QUIT

QU

Terminates break-in.

REMOTE

RU HOST *host*,*username*,PW *password*

Causes all subsequent RUNJOB or, in batch use, JOB commands to be passed to the specified host rather than to the host to which the terminal or batch entry station is connected. The setting may be overridden for individual RJ or JB commands by the inclusion of this command as a parameter. Before a TRANSFER command is issued a REMOTE command specifying the same host as the remote host in the transfer must be given. REMOTE without parameters obtains a report of the settings currently in effect.

RENAME

RN *file*,*file*

Changes the file name, or other part of the file description, of the first specified file to that specified second. May be used for magnetic tapes, when (*MT) must follow the first file name.

REPORT

RP *monitoraction*

Sets the level of monitoring information displayed on the terminal (see page 178.3 of the ICL Operating Systems GEORGE 3 and 4 manual). Default is FULLBUT,COMMANDS.

RETRIEVE

RV *file*,*file*,...

Causes the specified files (maximum of 23) to be restored if they have been placed off line. Other commands may be issued while the retrieval is taking place.

RETURN

RT *tape*

Releases the tape from the ownership of the user.

RUNJOB

RJ *jobname*,*username*,*file*,PARAM(*param*,*param*,...), JD(*scheduling*),OUTPUT(HOST *host*,*username*,PW *password*, PR *property*),REMOTE(HOST *host*,*username*,PW *password*)

Submits batch job held in *file*. *username* required only if LOGIN command has not been used. Parameters are passed to the job. See local documentation for details of *scheduling*. OUTPUT parameters determine the location at which any output from the job will appear, and REMOTE passes the job to a remote host (see the OUTPUT and REMOTE commands), defaults are the printer associated with the source of the command and the host to which the command is input. If the job is submitted to UMRCC, CP76 may replace JD to pass the job on to the CDC 7600; the scheduling parameters should then be appropriate for the 7600.

SCREENEDIT

SD *file*,*file*

Available on VDUs to call the on-screen editor at sites where this is provided. First specified file is the file to be edited, second the output file, default a file with the same name as but a generation number one higher than the edited file.

TRACE

TA *monitoraction*

Limits the information sent to the monitor file. See the ICL manual Operating Systems GEORGE 3 and 4 p 178.3 for an explanation of *monitoraction*.

TRANSFER

TF FROM *file*,HOST *host*,TO *file*,HOST *host*

Transfers first specified file from the first host to the second, where it is given the second file description. Default for either file description is the other file description, for either host is the host to which the command is input. When this command is issued, a REMOTE command specifying the same remote host must be in force.

TRAPCHECK

TC *file*,*username*

Obtains a list of the types of access the specified user (default the issuing user) has to the file.

TRAPGO

TG *file*,*username*,*mode*,*mode*,...

Removes restrictions on the specified user's access to the file owned by the issuing user. Modes are READ, WRITE, APPEND, EXECUTE, ERASE or ALL.

TRAPLIST

TL *file*

Lists all the access restrictions currently set on the file, which must belong to the user.

TRAPSTOP

TS *file*,*username*,*mode*,*mode*,...

Places restrictions on the access the specified user can have to the file. See TRAPGO.

WHATNWCOM

WN *commandtype*,HOST *host*

Obtains the number of the user's commands of the specified type (LF, JB, RJ, TFB, TFS, ALL) outstanding on the specified host (default all hosts).

WHATSTATE

WS ALL

Obtains a list of all the user's jobs on the current host that are yet to be completed.

N

W

UNIVERSITIES COMPUTER SERVICES

NWD 25

NETWORK GEORGE COMMANDS

Appendix 1

F D Sherwood
Keele
May 1979

This card summarises the main job control commands that are available on the ICL 1900 Series computers that are connected to the North West Universities Computer Network. The commands may be used at a MOP terminal connected directly to one of these machines, or at a Network terminal connected via a Network Interface Processor, or, where appropriate, in the job control files of batch jobs. If a Network terminal is used, connection to the required computer must first be established: press the CTRL and A keys of the terminal, then type

@CONNECT host

where host is the location of the computer to which connection is required, e.g. LIVERPOOL. If the host computer is available you will now be able to log in. If it is not available you will receive a message to that effect. To check the availability of the network host processors before attempting connection, use the command

@STATUS host

where host may be any location, or ALL for all hosts.

ABANDON

AB *jobname*,ME(*message*)

Abandons specified job, *message* appears in the monitor file.

ATTRIBUTE

AU PR *property*

Any output produced by LISTFILE, other than to the terminal, will appear at the outstation or particular device defined by *property* on the host being used. The default property is that associated with the terminal or batch input device from which the command originated, but for output from remote hosts see OUTPUT.

CANCEL

CC *command parameters*

Cancels a previous ATTRIBUTE, QUIET, RETRIEVE or TRANSFER command, specified in the same form as it was originally given.

CONTINUE

CU

Resumes job after break-in.

COPY

CY *file*,*file*

Contents of first file specified are copied into the second file.

@DISCONNECT

@D

Used only at Network terminals to disconnect from the host after LOGOUT. Alternatively @C (or @CONNECT) may be used to connect to a different host.

EDIT

ED *file*,*file*,*file*

Calls the GEORGE editor. The first file is the file to be edited, the second the file to receive the edited version (default a file with the same name as but generation number one higher than the edited file), the third the file containing the editing instructions (default the terminal or the following records in a batch job).

ENDJOB

EJ *monitoraction*,RETAIN(*file*)

Used only in batch jobs to terminate. If *file* is specified, the monitor file will be held there as well as being printed. *monitoraction* determines the contents of the monitor file, see p 178.3 of the ICL manual Operating Systems GEORGE 3 and 4.

ENTER

EN *n*,PARAM(*param*,*param*,...)

Causes loaded object program to be entered at entry point *n*. Parameters are passed to the program.

ERASE

ER *file*,*file*,...

Deletes all the specified files (upto 23 may be specified).

EXIT

EX

Used in macros to terminate the obeying of commands from the macro file.

GET

GE *tapename*(*MT)

Obtains and names a magnetic tape for the user.

GOTO

GO *label*

Used in macros to jump to a labelled command.

IF

IF *condition*,(*command*)ELSE(*command*)

First command is obeyed only if condition is true, second command if false. ELSE(*command*) may be omitted, when the next command will be obeyed if the condition is false. See the ICL manual Operating Systems GEORGE 3 and 4 p 284 for details of conditions.

INPUT

IN *file*,T*terminator*

Copies the following text, up to the terminator, into the file. If T is replaced by S, the terminator is also copied. If T*terminator* is omitted, S**** is implied.

JOB

JB *jobname*,*username*,T*terminator*,PARAM(*param*,*param*,...),JD(*scheduling*),REMOTE(HOST *host*,*username*,PW *password*),OUTPUT(HOST *host*,*username*,PW *password*,PR *property*)

Introduces batch job; may be used at a terminal only when the user is not logged in. The records following, up to the terminator (default ****), will be submitted as a job. The parameters after PARAM, if any, will be substituted for %A, %B etc in the job description. See local documentation for details of scheduling parameters. REMOTE causes the job to be transmitted to the specified host (see the REMOTE command), and OUTPUT causes output to be sent to a remote site (see the OUTPUT command). If the job is submitted to UMRCC, CP76 may replace JD to pass the job on to the CDC 7600; scheduling parameters should then be appropriate for a 7600 job.

LISTDIR

LD *directoryname*,*level*,*LP,PR *property*

Lists contents of directory with HIGH or LOW level of detail, on the terminal (or in the monitor file) if *LP and PR omitted.
Note: Parameters are positional, so commas must always be present.

LISTFILE

LF *file*,*device*,FROM*l*,TO*m*,LINES*n*,NUMBER,PR *property*,OUTPUT(HOST *host*,*username*,PW *password*,PR *property*),LOCAL

Causes output of *file* on *device* (*LP etc, default the terminal). *l* is a line number defining the start of a portion of the file to be listed (default start of file), *m* is the end line number or *n* the number of lines (default the rest of the file). NUMBER obtains line numbers. The first PR sends output to the specified outstation or particular device on the host being used, otherwise the property in the OUTPUT parameter or a previous OUTPUT command will be used (see the description of the OUTPUT command). LOCAL may be used to temporarily override the OUTPUT command currently in effect and produces output at the site to which you are connected.

LOAD

LO *file*,CORE *n*

Loads binary program from the file, and requests *n* words of store.

LOADENTER

LE *file*,*n*,PARAM(*param*,*param*,...)

Loads a binary program from *file* and enters it at entry point *n*. Parameters are passed to the program.

LOGIN

LN *jobname*,*username*,JD(*scheduling*),REMOTE(HOST *host*,*username*,PW *password*),OUTPUT(HOST *host*,*username*,PW *password*,PR *property*)

Used to log in at a terminal. See local documentation for details of job name conventions and scheduling parameters. REMOTE and OUTPUT have the same effects as separate REMOTE and OUTPUT commands.

LOGOUT

LT *monitoraction*,RETAIN(*file*)

Used at a terminal to log out. Parameters are as in ENDJOB.

MACDEF

MD *file*,T*terminator*

Stores the following records, upto the terminator, in the file as a macro. Default terminator is ENDM.

MAIL

ML LIST,KEEP

Obtains display of messages sent to the user. Messages are then deleted from store unless KEEP is specified.

ML username,message

Sends message to the specified user.

MSGNETWORK

MG NEST,*LP

Obtains status information on all network hosts.

NCWABANDON

NB *command*,HOST *host*,MULTI

Causes commands referring to the specified host that have not yet been executed to be abandoned. Commands should be specified as LF *filename*, RJ *jobname* etc. If only the command name is given, any such commands will be abandoned. If MULTI is omitted and a parameter is given with the command name, only the longest standing occurrence of the particular command will be abandoned. TRANSFER commands should be specified as TFB for bring or TFS for send. *command* may be ALL for all commands.

NCWLIST

NL *commandtype*,HOST *host*,LIST

Obtains a list of the user's outstanding commands of the type specified (LF, JB, RJ, TFB, TFS, ALL) relating to the host (default all hosts), on a line printer if LIST is present.

NCWSTAT

NS

Obtains the status of the network command well, the number of outstanding commands, and the number of commands the well can hold.

NETWORKING IN THE MIDLANDS UNIVERSITIES

The six Midlands Universities, Aston, Birmingham, Leicester, Loughborough, Nottingham and Warwick are developing a packet switched network, MIDNET to link the computing facilities of each University, enhance existing access to computing facilities at the University of Manchester Regional Computing Centre (UMRCC) and provide access to the Post Office Packet Switched Service (PSS).

The existing situation is shown in Figure 1. There are no direct links between sites and access to UMRCC uses 7020 links via the JANUS and SWAN facilities at Nottingham and Birmingham. Job transfer to UMRCC is "cards in, paper out" except at Nottingham and Birmingham.

The MIDNET topology is shown in Figure 2. Each site has a local MIDNET switching node. These nodes are initially connected in a ring as shown giving two possible routes between any two sites. MIDNET will provide full access to the computing facilities at each site i.e. file transfer, job transfer, interactive working. Two connections to PSS are proposed, at Birmingham and Nottingham, to provide access for all MIDNET sites. The connection to UMRCC is enhanced to become a true network connection.

Each MIDNET node is a SYSTIME 5100 (PDP11 11/34) with 128K bytes of MOS store, 2 RK05 compatible disk drives, DUP11 synchronous interfaces and DL11 asynchronous interfaces. The internode links run at 4.8K bits/sec.

The mainframe computing facilities at each site are summarised below.

Computing Facilities at Each Site

Aston	1904S
Birmingham	1906A, DEC2050
Leicester	Cyber 73, local campus network
Loughborough	1904A, Twin Prime 400
Nottingham	1904S, 1906A, local campus network
Warwick	Burroughs B6700, local campus network

Structure of MIDNET

MIDNET conforms to the level structured approach. Figure 3 illustrates the components of MIDNET, and shows the level at which each component operates, and the protocols employed at each level. The software is modular and, in general, a module represents a MIDNET component, operating at the required level with well defined interfaces between itself and the modules at adjacent levels.

A PROCESS provides access to the network for a user or access to a service from the network. In general, one multithreaded task or program will implement several identical processes, each process supporting one call.

A PIN (Process Interface Node) provides a transport service interface to the network for one or more processes. It represents the end point of a virtual call at level 3 and controls the network dependent features of the call (e.g. flow control, packet formats).

A NETEX is a packet switching exchange. It has a functional role which involves establishing the initial route a call will use, maintaining its links and dealing with link failures, switching subsequent packets, and clearing calls. It also has a service role which involves maintaining information about current traffic and link states, and providing an interface with the outside world to allow inspection and control operations to be performed. This control interface is currently accessed via the console terminal on the node. An additional facility provided by the NETEX allows communication between the console terminals on each node by means of what is best described as a "permanent virtual ring". This facility is independent from the packet switching function of the NETEX, and is proving a valuable aid to testing between sites.

Link modules are treated separately from the logical components since they need not appear in a logical diagram of the network (i.e. a diagram which does not indicate the underlying hardware connections). They are inserted where required when moving from the logical diagram to the physical hardware diagram.

Although X25 level 1 defines an electricalinterface we find it convenient to talk in terms of a level 1 module which controls that interface and provides a software interface for a level 2 module. The level 1 module on the MIDNET nodes controls a DUP11 interface which provides CRC generation and error detection, flag generation and bit stuffing in hardware.

Relationship Between Logical Components & Physical Components

A network is normally viewed as a black box, provided by the network administration, with well defined interfaces to which subscribers to the network may connect their equipment. There is a strict separation between "network" equipment and "subscriber" equipment (cf. DCE & DTE). Since both the administration and subscribers belong to the same body, in MIDNET a more flexible approach is possible. The logical components of MIDNET may be sited as required within the available hardware. Figure 4 shows some of the possibilities with regard to the connection of a mainframe to a MIDNET node.

Processes may be implemented in the node and the mainframe may support a NETEX if required. Thus the network - subscriber boundary need not coincide with an interface between separate equipment.

Standard DCE - DTE interfaces will also be supported by MIDNET for connections to PSS and UMRCC, and to allow replacement mainframes with supplied X25 interfaces to be readily connected.

Protocols in MIDNET

Level 1

The internode links comply with level 1 of X25, implemented on a V24 modem interface. As stated previously, a level 1 Device module (software) handles the level 1 interface and some of the lower level 2 functions (e.g. CRC generation, flag generation, etc.) in order to provide a software interface for level 2 modules.

The node to mainframe links employ native protocols supported by the local mainframes. The peculiarities of such a link (7 bits "wide", half duplex, polling, etc) are hidden by the level 1 software (using data compression, etc.) such that the link appears to be 8 bits "wide" and full duplex to the level 2 modules. The level 2 - level 1 software interface is standard throughout MIDNET, allowing the same level 2 module to be employed for internode links and node to mainframe links. When a mainframe is replaced by one which offers an X25 interface, the native level 1 module may be easily replaced by a standard X25 module without changing the higher level software. Figure 5 illustrates the simplicity which can be achieved in connecting a mainframe to MIDNET.

Level 2

The links to UMRCC and PSS will employ standard LAPB. The internode links employ a slightly modified procedure called "LAPM". The reason for this is that X25 level 2 defines an interface procedure between a network and subscribers equipment, whereas the internode links are inside the network and require a more symmetrical procedure. The level 2 COMMS module for the node will be easily configurable to support LAPB or LAPM.

Level 3

MIDNET employs a call level protocol called X25M which is identical to X25 level 3 in all aspects except the following:-

(i) Procedures have end-to-end significance - X25M is a protocol rather than an interface specification.

(ii) REJ and RESTART packets do not occur in X25M.

(iii) The Call Request and Call Accept packets always have a 128 byte (maximum) Call User Data field (cf. PSS Fast Select format).

(iv) The Call Accept packet contains address and facility fields (cf. PSS Extended Format).

(v) Call clearing is not destructive to any in-flight data packets.

Level 4

The interface provided to processes is fully compatible with the PO SG3 Transport Service proposals.

The implementation of the transport service has been simplified in the following ways:-

(i) A one to one correspondence between an X25M packet and a transport service primitive exists.

(ii) A limit of 128 bytes is imposed on dormant address information.

(iii) PUSH is implicit with a complete packet sequence (an explicit PUSH may be represented by a null packet sequence).

(iv) The unit of data may be selected by the process. The Physical process - PIN interface provides a local fragmentation scheme for transfer of data between the process and the transport service, if required.

High Level Protocols

FTP

The "blue book" definition will be used. Each site has defined its mapping between the conceptual filestore and its real filestore and defined the attributes it will support at the Q end. Rather than define a MIDNET subset, each site will support as many attributes as it can.

The P end will initially provide a "batch mode" user interface which allows a user to queue a request for file transfer which is subsequently executed. More "user friendly" interfaces will be developed subsequently. The P end will allow the user to specify any attribute with his request (i.e. limitations to the range of attributes supported occur at the Q end, not the P end).

JTP

Until a standard arrives, MIDNET will employ interim measures built around FTP, without modifying FTP. The interim measure relies on sites being able to recognise the source site when a "Take Job Input" occurs, store that information while the job awaits execution, and put job output in a known

place where it can be recognised by a subsequent "Give Job Output" request. Interactive facilities will be developed to provide a separate means of interrogation and control.

Interactive Working

It is intended to use Triple X (X3,X28,X29). However, difficulties arise when applying Triple X to terminals attached to mainframes and online services on mainframes, which are message buffering and employ strict conversational control over the terminal. An in-house solution could be provided but would impede open systems working, so we are extremely interested in forthcoming proposals from the Character Terminals Working Group.

PSS Connections

Figure 6 shows the logical diagram of the gateway between MIDNET and PSS. The Transport Station Node will operate as described in the draft Transport Service specification (yellow book). The gateway will also perform Access Control and Accounting but details of these functions have not been fully studied yet.

Progress to Date

Most of the node software for the initial service is complete. This includes the NETEX, COMMS, DEVICE, PIN modules along with the supporting software for control and monitoring. Test modules have been produced for testing individual components in a "single shot" manner. A test process is being produced to facilitate testing of throughput and response under varying loading conditions and a simple PAD has been produced to access the test process from a node terminal. A simple FTP process (P and Q) for the node is nearing completion.

Most sites now have a working node to mainframe link and are concentrating on implementation of FTP processes in the mainframe.

An experimental service will start at the end of November 1979. This will initially operate for a few hours per day and provide file transfer facilities. Most sites will have their FTP processes ready for the start.

FIGURE 1 EXISTING MIDLANDS CONNECTIONS

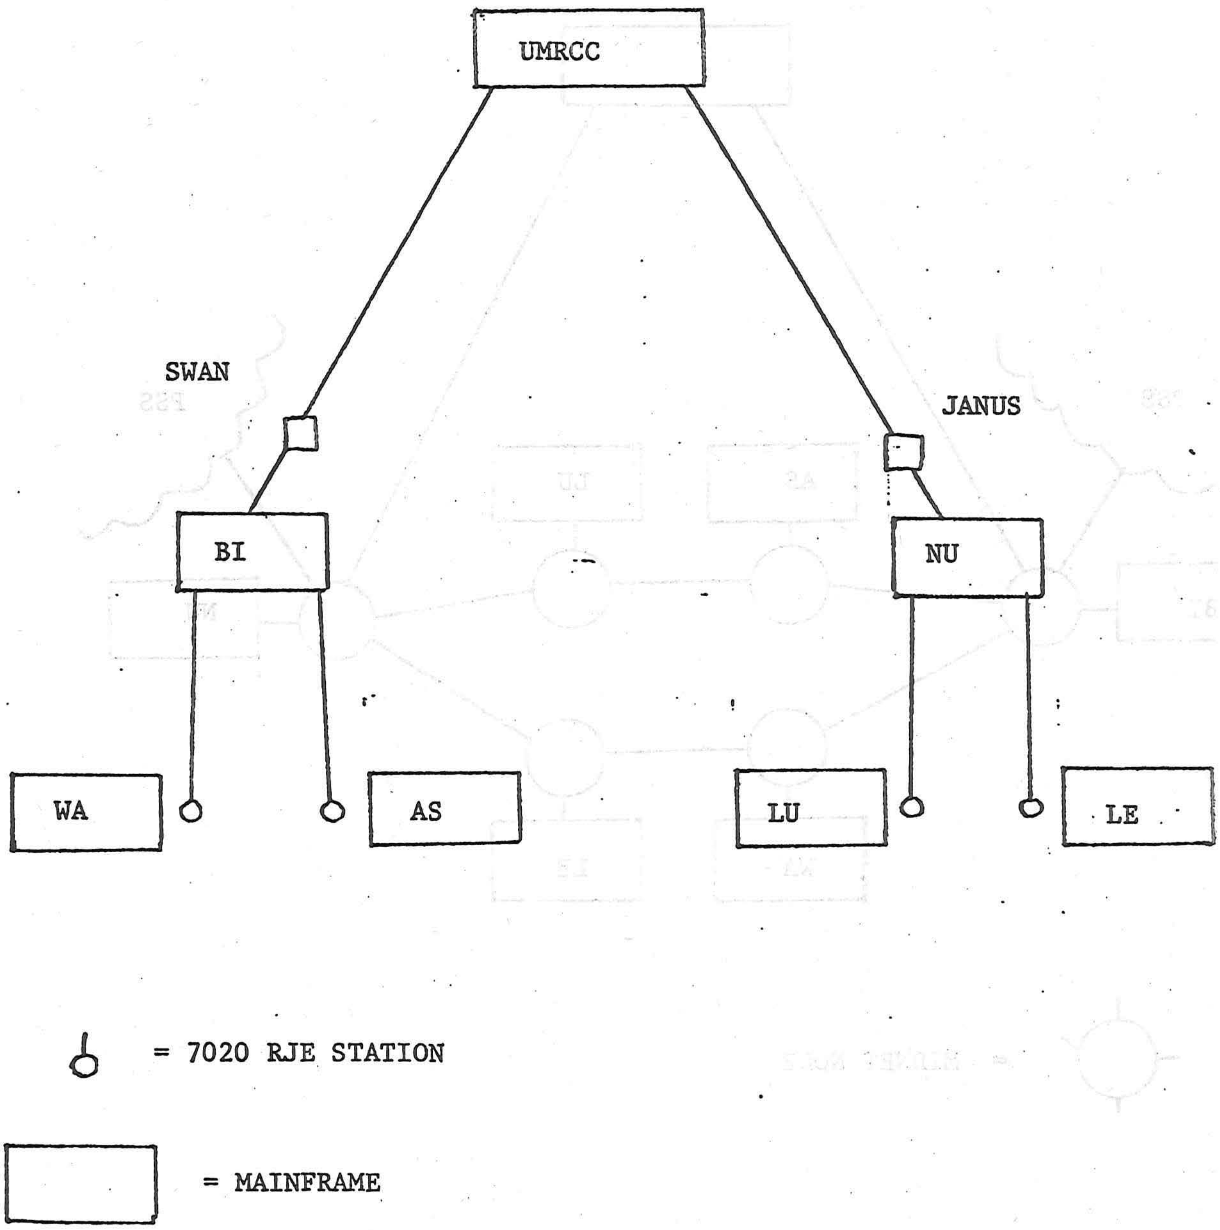

FIGURE 2 MIDNET TOPOLOGY

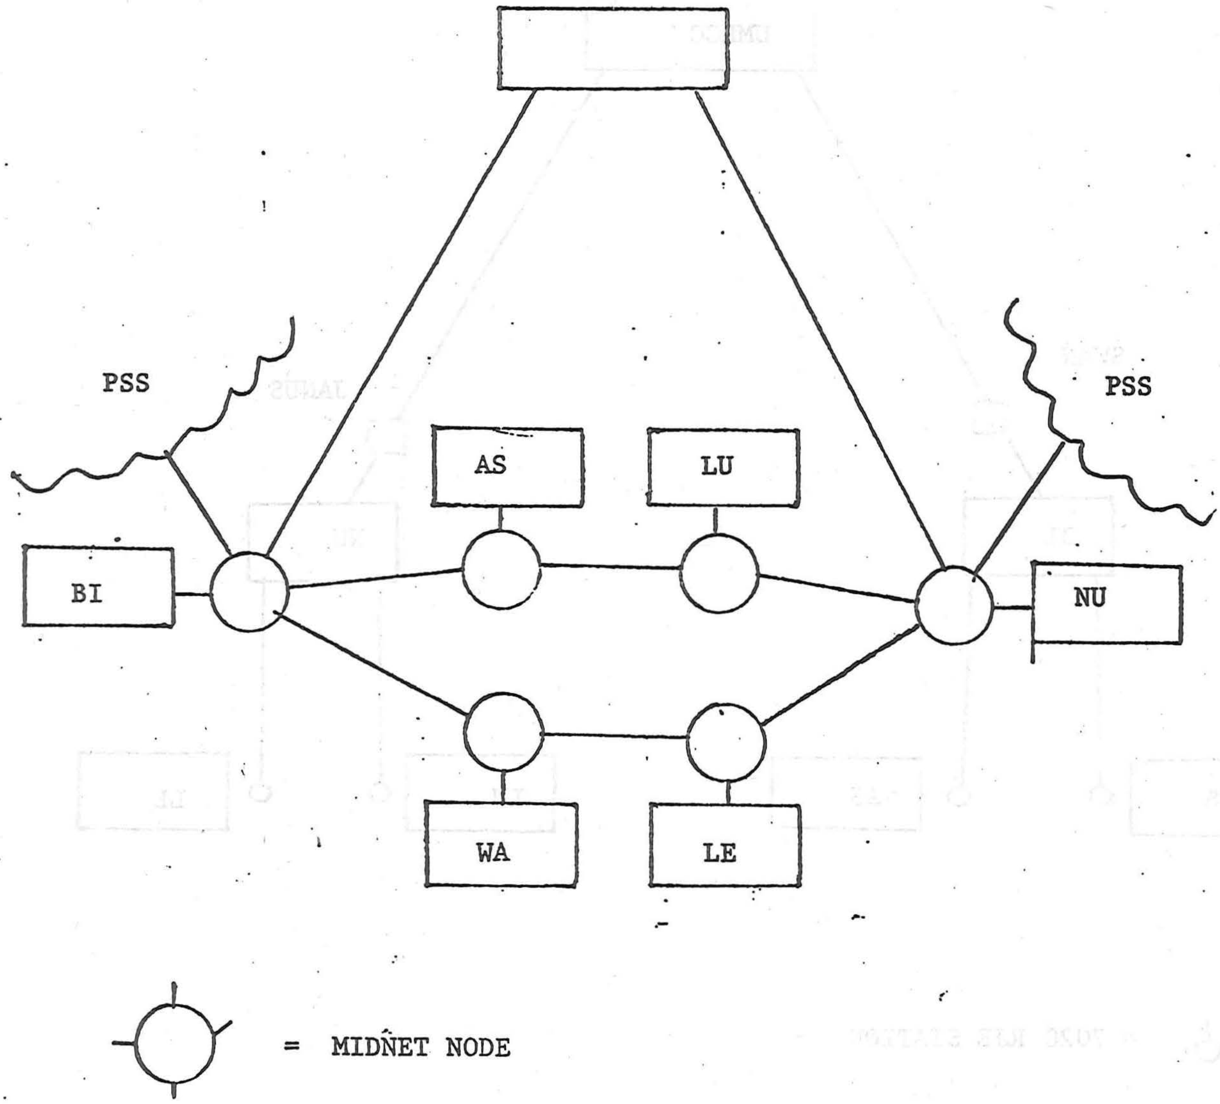

FIGURE 3 LOGICAL COMPONENTS

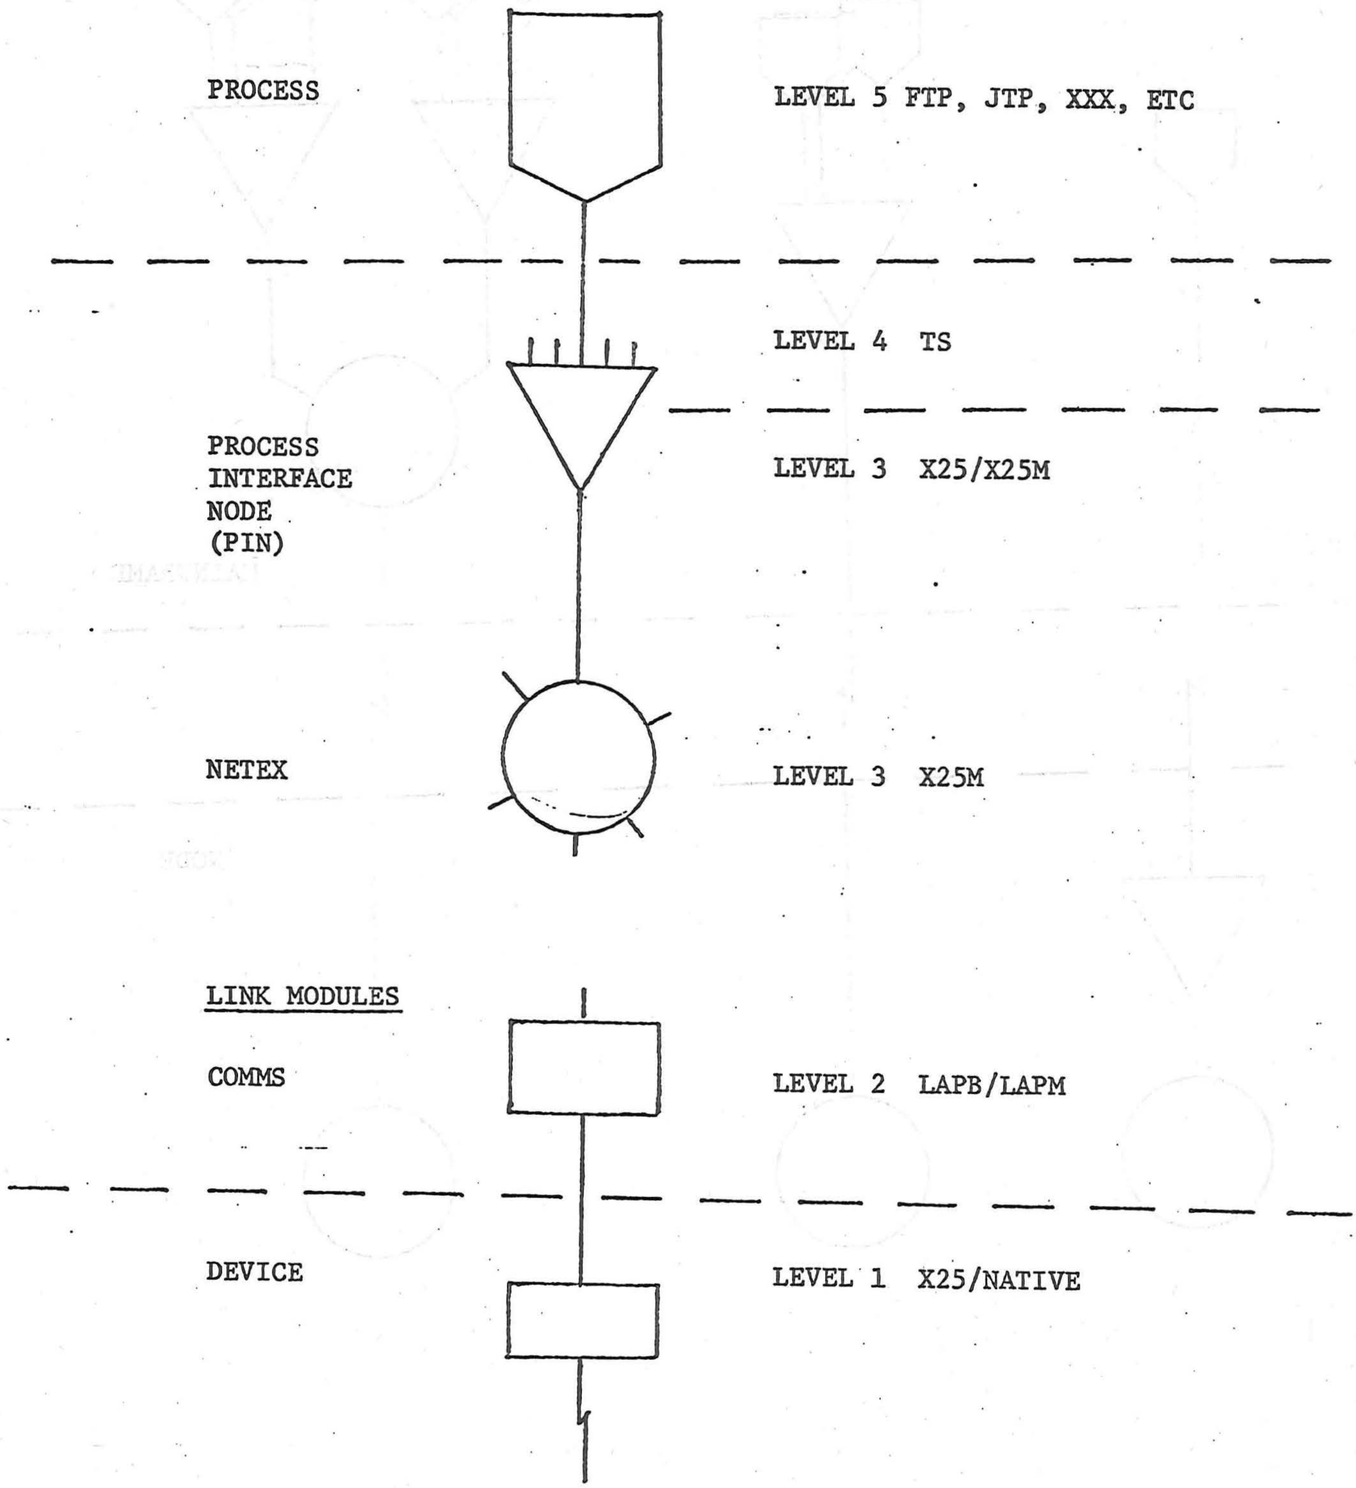

FIGURE 4 LOGICAL CONNECTIONS

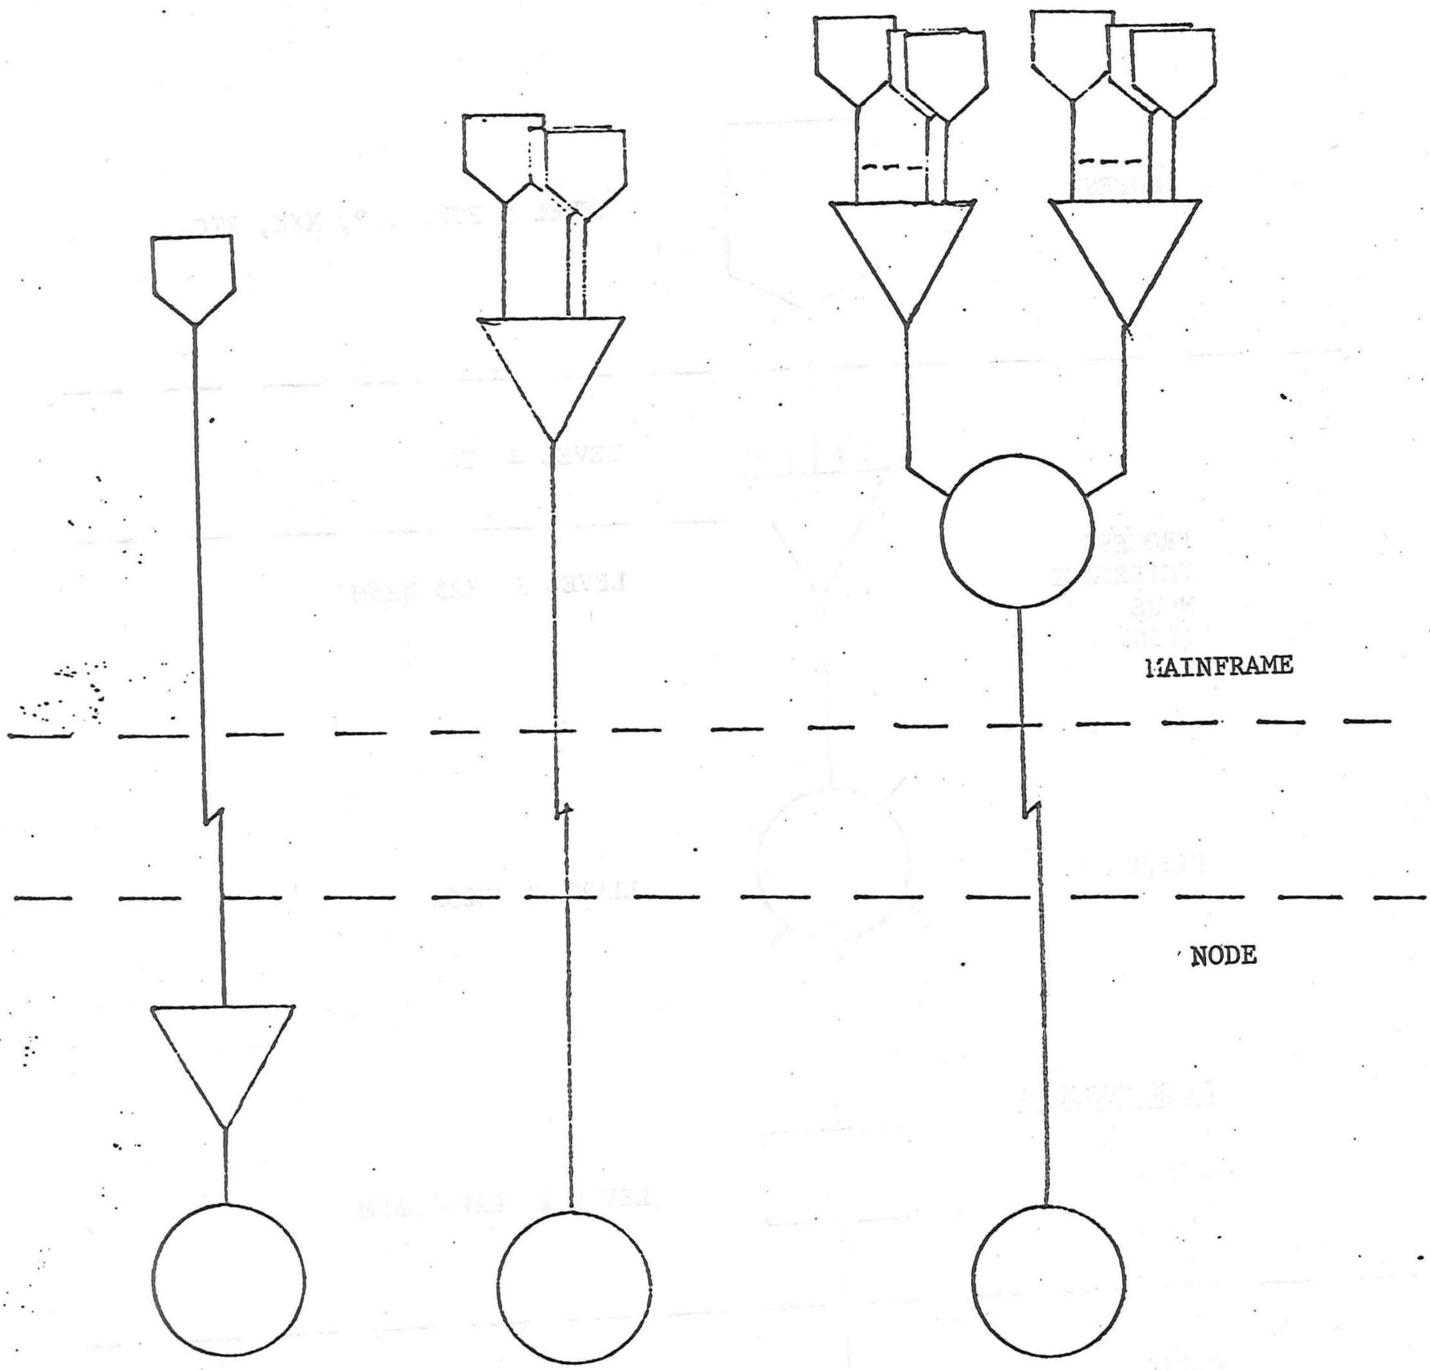

FIGURE 5 PHYSICAL CONNECTIONS

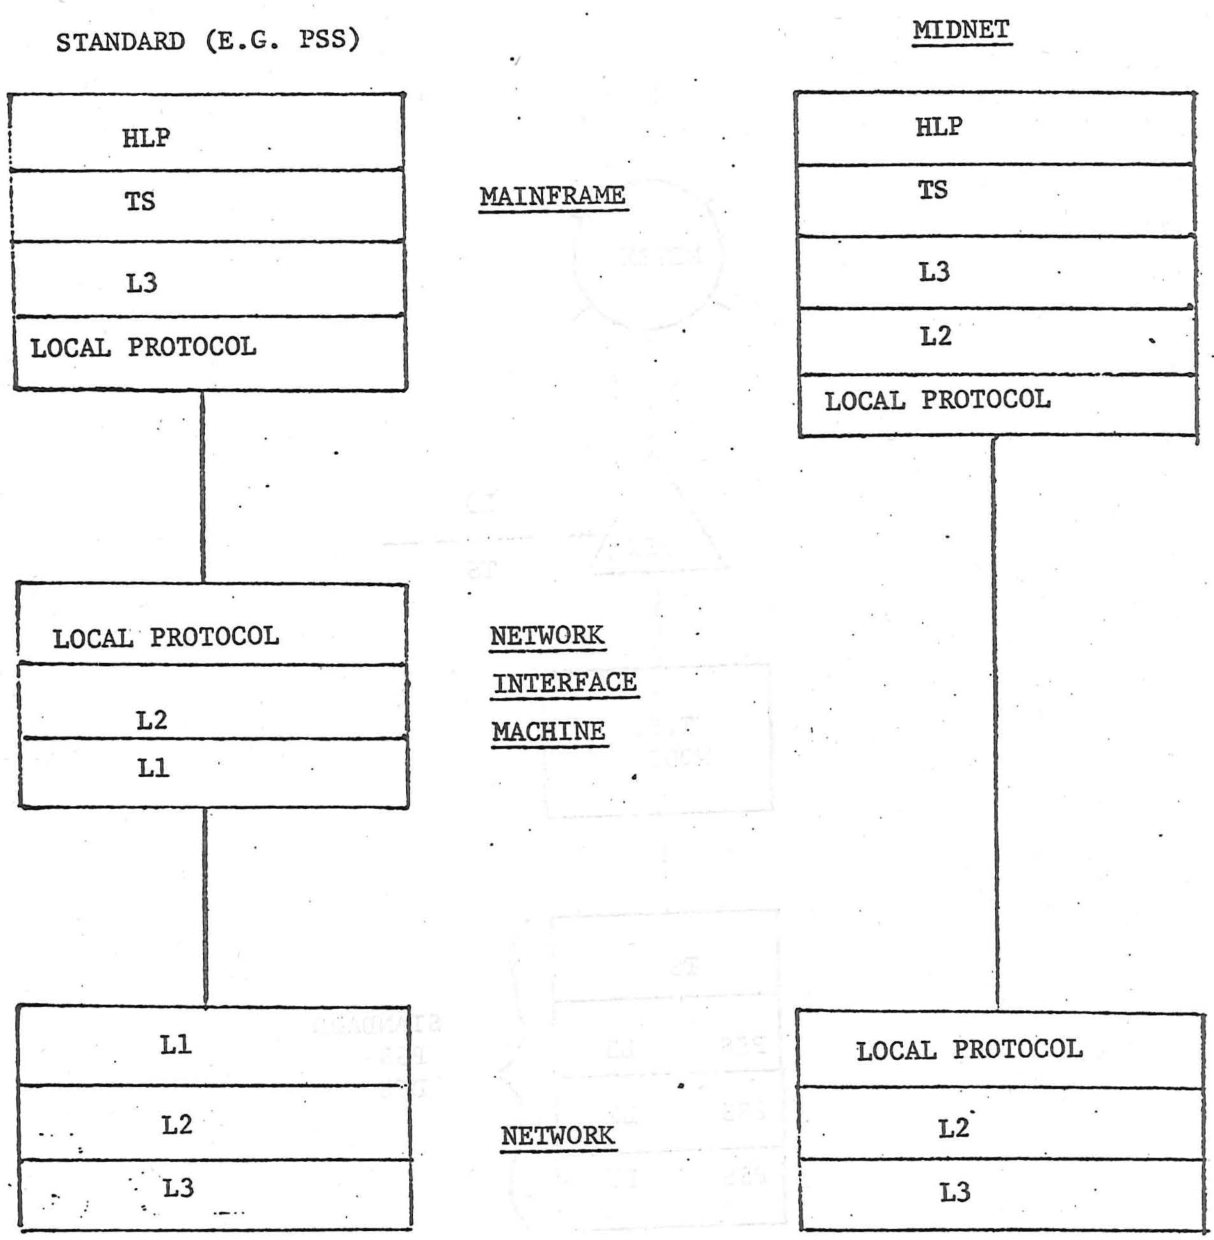

FIGURE 6 PSS CONNECTIONS

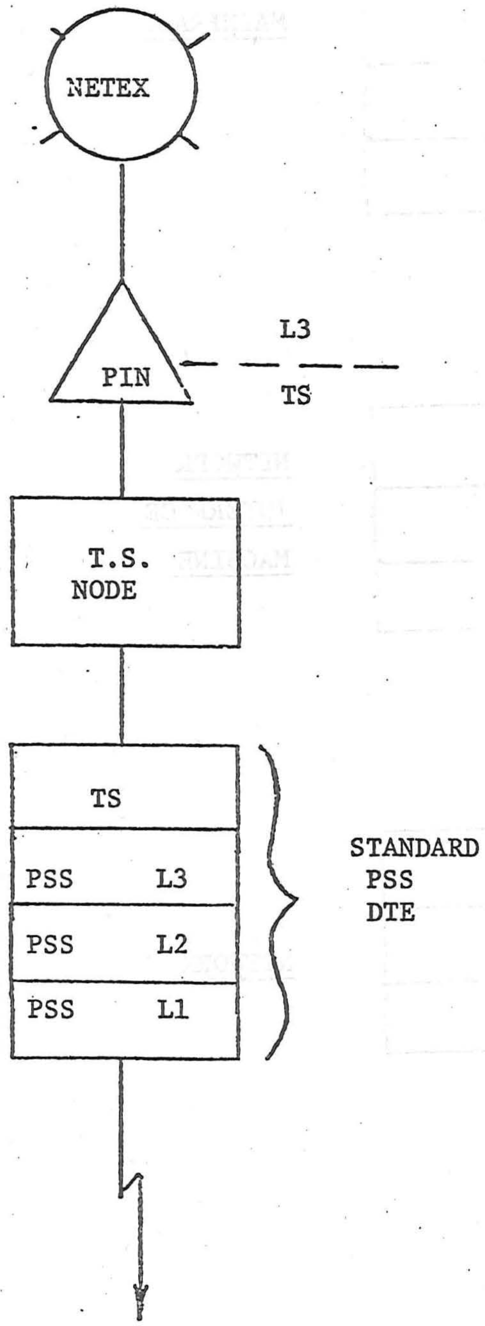

FIGURE 7 A TYPICAL MIDNET NODE

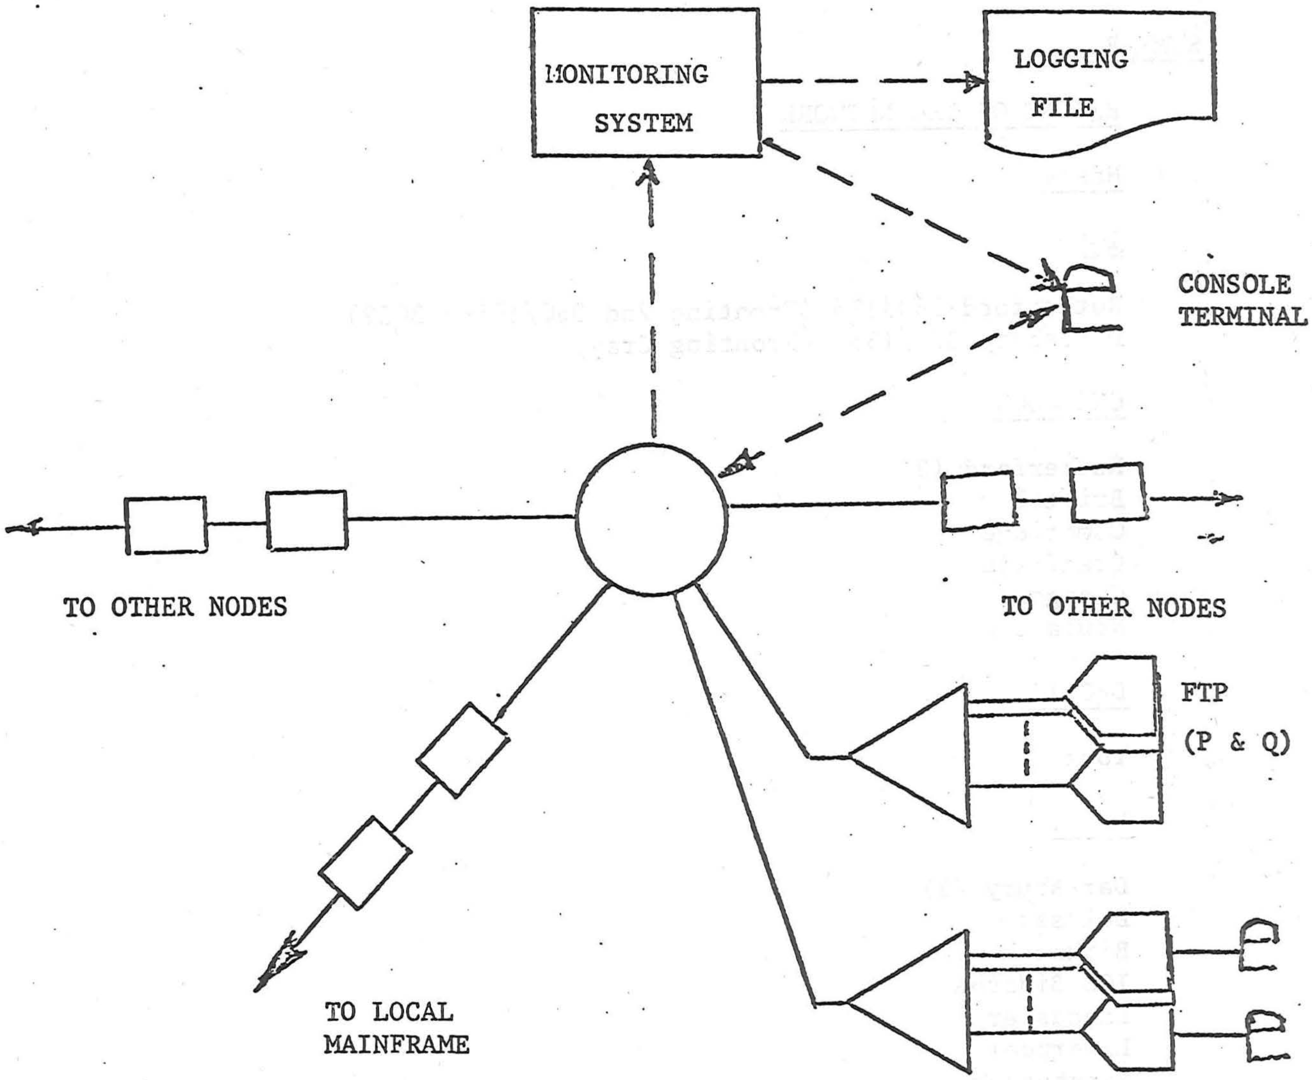

RUTHERFORD & APPLETON LABORATORIES
COMPUTING DIVISION

CURRENT STATE OF THE SRC X25/EPSS NETWORK

J W Burren

15 October 1979

SUMMARY

1. EXTENT OF THE NETWORK

1.1 Hosts

IBM

Rutherford 360/195 (Fronting 2nd 360/195 + 3032)
Daresbury 370/165 (Fronting Cray)

GEC 4000

Rutherford (2)
Bristol
Cambridge
Cranfield
Glasgow
Newcastle

DEC 10

York

PDP 11

Daresbury (2)
Belfast
Birkbeck
IOS Bidston
Lancaster
Liverpool
Manchester
Sheffield
SMBA Oban
York

GEC 2050

Rutherford (2)
Bristol
Southampton

1.2 Exchanges

Rutherford	GEC 4000
Daresbury	PDP 11's

1.3 Configuration

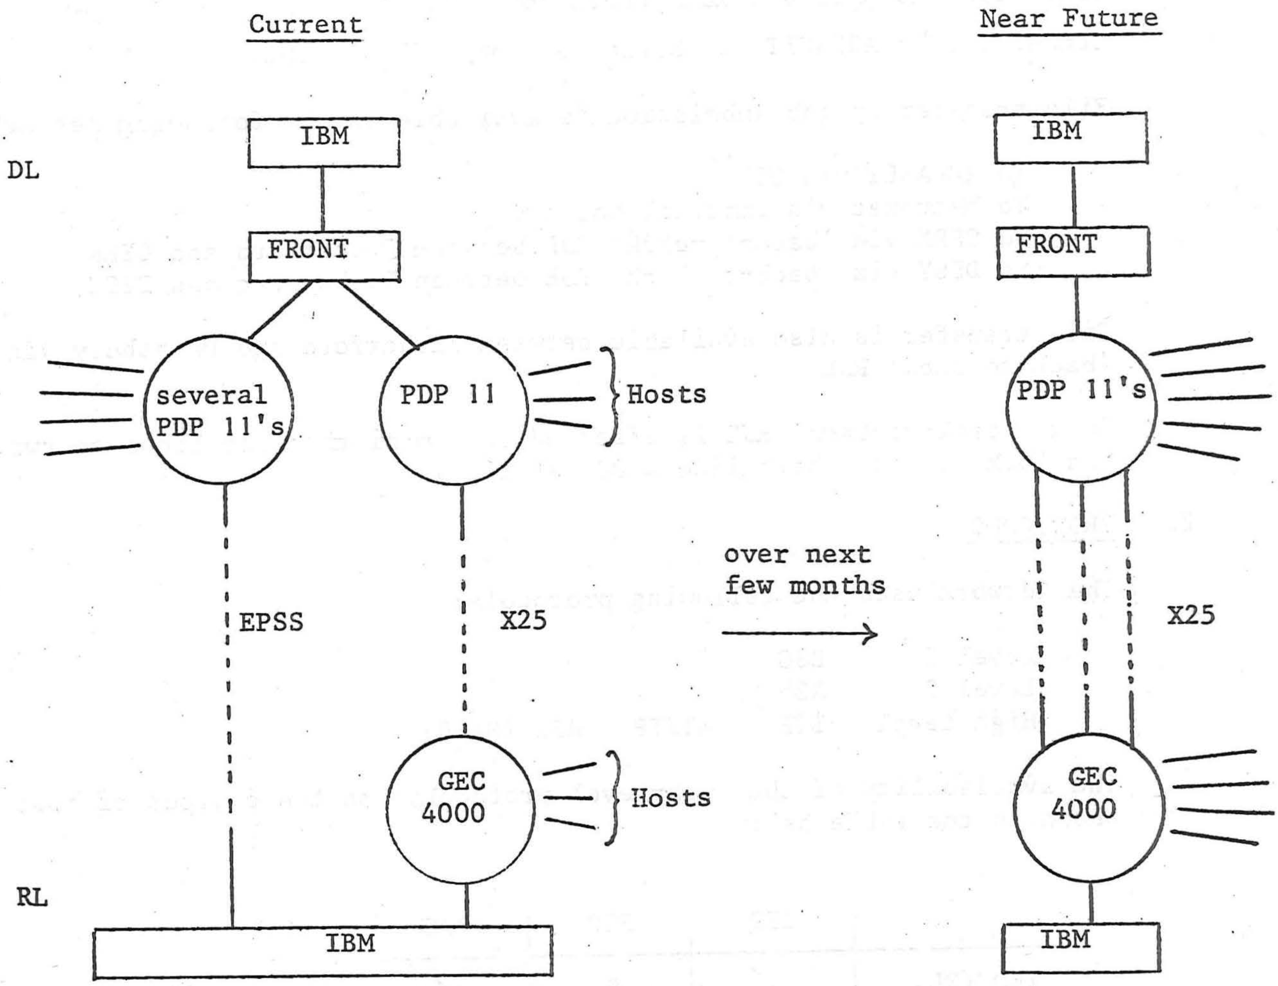

1.4 Traffic in Network

Unfortunately there is currently no systematic way of collecting traffic statistics for the network. Some known traffic measurements are:-

Typical weekly traffic into 360/195

5000 calls/week (4,400 terminal, 600 RJE)
1 million packets (500,000 terminal, 500,000 RJE)
100 million bytes (25M terminal, 75M RJE)

Traffic through RL exchange (GEC 4000)

Maximum hourly average for previous week
1000 bytes/sec averaged during 1 hour
150 calls made during this hour
Maximum concurrent calls 20.

1.5 Gateways to other Networks

Terminal calls may be made between the following networks and the SRC X25 network with the IBM mainframes acting as gateways.

Daresbury internal network:

NSF Control computers (Interdata)
NSF Data Collection computers (GEC 4000)
SRS Control computers (Interdata)
SRS Data Collection Computers (PDP 11)
Terminal concentrator (PDP11)

NSF = Nuclear Structure Facility

SRS = Synchrotron Radiation Source

Rutherford RJE 'star' network and terminal network.

Connection to EPSS via Rutherford 360/195.

Connection to ARPANET via Rutherford 360/195 and UCL.

File transfer by job submission is available to the following networks:

To ARPANET via UCL
To Metronet via Imperial College
To CERN via 'back-to-back' RJE between Rutherford and CERN
To DESY via 'back-to-back' RJE between Rutherford and DESY.

File transfer is also available between Rutherford and Daresbury via 'back-to-back' RJE.

(Note: 'back-to-back' RJE is effected on a mini that has lines to two hosts and looks to each host like a RJE station).

2. PROTOCOLS

The network uses the following protocols:

Level 2	BSC		
Level 3	X25		
High Level	ITP	NIFTP	RJE (HASP)

The availability of the high-level protocols on the 6 types of host is shown in the table below.

	ITP	FTP	RJE
IBM (RL)	✓	S	✓
IBM (DL)	✓	P	✓
GEC 4000	✓	✓	✓
DEC 10	✓	P	N
PDP 11	✓	M	✓
GEC 2050	✓	N	✓

S = Soon
P = Planned
N = No
M = Minimal subset

3. GENERAL REMARKS

3.1 Most of the network works very well, but it is quite hard to keep all of it working well all the time.

3.2 In addition to a Network Development Committee (currently at meeting 28), we now have a Network Operations Committee (currently at meeting 3). The first priority request from the Operations Committee to the Development Committee is for a Status machine to give status information on all components of the network, to collect statistics, to give broadcast facilities and login messages, etc.

3.3 The ITP protocol is a success and allows good terminal access to hosts having widely different terminal handling systems.

4. FUTURE DEVELOPMENTS

4.1 Operational considerations will dominate future development.

4.2 Most important immediate task is to get rid of EPSS hangovers.

4.3 Next most important task is development of a STATUS machine.

4.4 It is the intention of the Development Committee to move as close as possible to the PSS definition of level 3. A first look at the Post Office specification has not shown any 'stoppers'.

4.5 PSS level 2 will be introduced as an option.

TRIPLE-X AND THE TRANSPORT SERVICE

PETER HIGGINSON

University College, London

TRIPLE-X AND THE TRANSPORT SERVICE: SUMMARY OF A PRESENTATION AT NETWORKSHOP 5

by P.L. Higginson
University College London

1. General

This note is a summary of a presentation on the work of the Character Terminal Working Group of the PSS User Forum Study Group 3. The summary is short because the full report will be available shortly and will be distributed to the Networkshop attendees. Further copies of the full report can either be obtained from myself, as chairman of the Working Group, or from Post Office PSS Marketing.

2. Effects of a Network on the Host/Terminal Interface

Terminals local to a computer tend to be handled in different ways by different types of computers and, although some form of terminal handler is often present, its functions vary. Use of a network introduces:

- a greater variety of terminals,
- a delay in characters arriving at the host, and
- a preference for use of messages on the network rather than single characters (for cost and other reasons).

Some form of terminal handling is therefore required in the network for character terminals. Also, in the near future, there will be terminals with packet interfaces to the network, and these will need to use the same protocols.

3. Functions of the Packet Assembler and Disassembler (PAD)

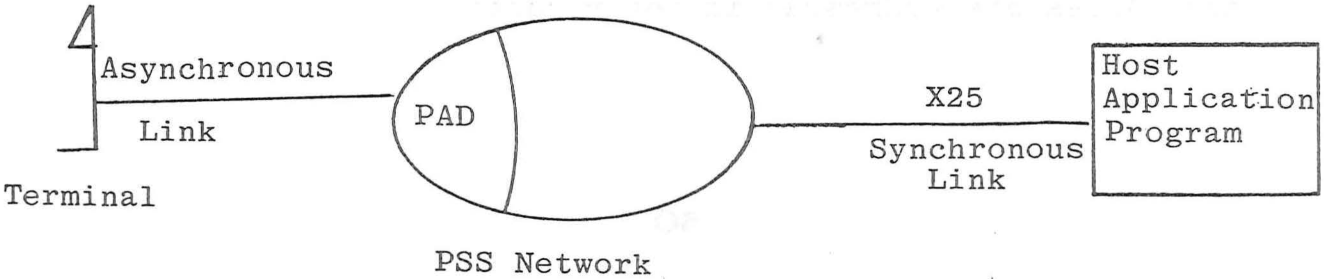

The group has considered the functions of the PAD and found the most important to be:

- Setup and control of the connection to the host
- Echoing of input as typed (for full duplex terminals)
- Editing of input
- Transmission of input to host
- Formatting of output from host
- Interleaving of echo with output

and two optional facilities which PADs may provide:

- Possibility of Transparent Output
- Use of auxiliary devices or files.

Ease of use is of considerable importance and therefore single keys should be provided for the editing functions (delete character, delete line and retype line), for Transmit, for "New line and Transmit" and for escaping into the PAD command mode. In addition a method must be provided to input all codes. On output the PAD should insert any necessary delays (or padding), should format according to page and line lengths, by folding long lines and having the capability to wait at end of page, and should generate correct parity if necessary.

4. Main Components of the Group's Report

The aim of the group's report is to propose solutions to the difficulties which it is felt will be encountered by implementors of X3, X28 and X29 on PSS. The difficulties fall into two categories; firstly, the problem of ensuring that different implementations are compatible with each other, and secondly the practical question of how far X3, X28 and X29 are adequate for the task of handling terminals across a network.

Three main issues are addressed in the report:

- In the interests of the interworking it is recommended that a single method of handling terminals connected across PSS be used. In conjunction with a special PAD terminal profile defined later in the report, this means that the PAD message protocol defined in X29 is kept to a bare minimum. It is hoped that this proposal will meet the widest possible acceptance amongst users of PSS, leading to a broad degree of compatibility.

- It is felt that some operational difficulties will arise in the use of X3/X28/X29 as they stand and so the report points out directions in which it is felt that improvements could be made, for example by proposing extra PAD parameters to which a host may have access, and by proposing ways in which X28 could be enhanced to provide a better user interface.

- Finally X3/X28/X29 are defined as using an X25 network. The group feels that this will lead to problems both in the short, and in the long term, when terminals may be connected through non-X25 networks. It therefore recommends the Transport Service as a more suitable bearer than a raw X25 network for the end-to-end functions of controlling remote terminals.

Although they are presented together, these proposals are functionally independent. Indeed, care has been taken to ensure that different levels of implementation of the report's proposals will remain compatible with each other, since it is vital, for example, that terminals connected to a standard PSS PAD should be able to access any host regardless of its implementation, and that terminals connected to private PADs or networks should be able to access standard host implementations of X29 wherever they may be.

We have chosen the name TS29 for the protocol we have defined, which includes a subset of X3/X28/X29, uses the Transport Service and contains some enhancements to X3/X28/X29. The enhancements could also be used with X3/X28/X29 directly.

5. <u>The Problems of Local Networks</u>

The problems of connecting local networks to PSS have been carefully studied from the point of view of terminal interworking. In particular, many current local networks or host terminal handling systems could not be mapped onto X3/X28/X29, and the TS29 and X29sub protocols are designed to make this connection easier. The use of Transport Service addressing capabilities is very necessary for local network to local network connection via PSS and to simplify the interface to the host applications. We foresee the necessity to provide conversion between the protocols of X28 network terminals and TS29 either in the gateway to a local network or in individual hosts.

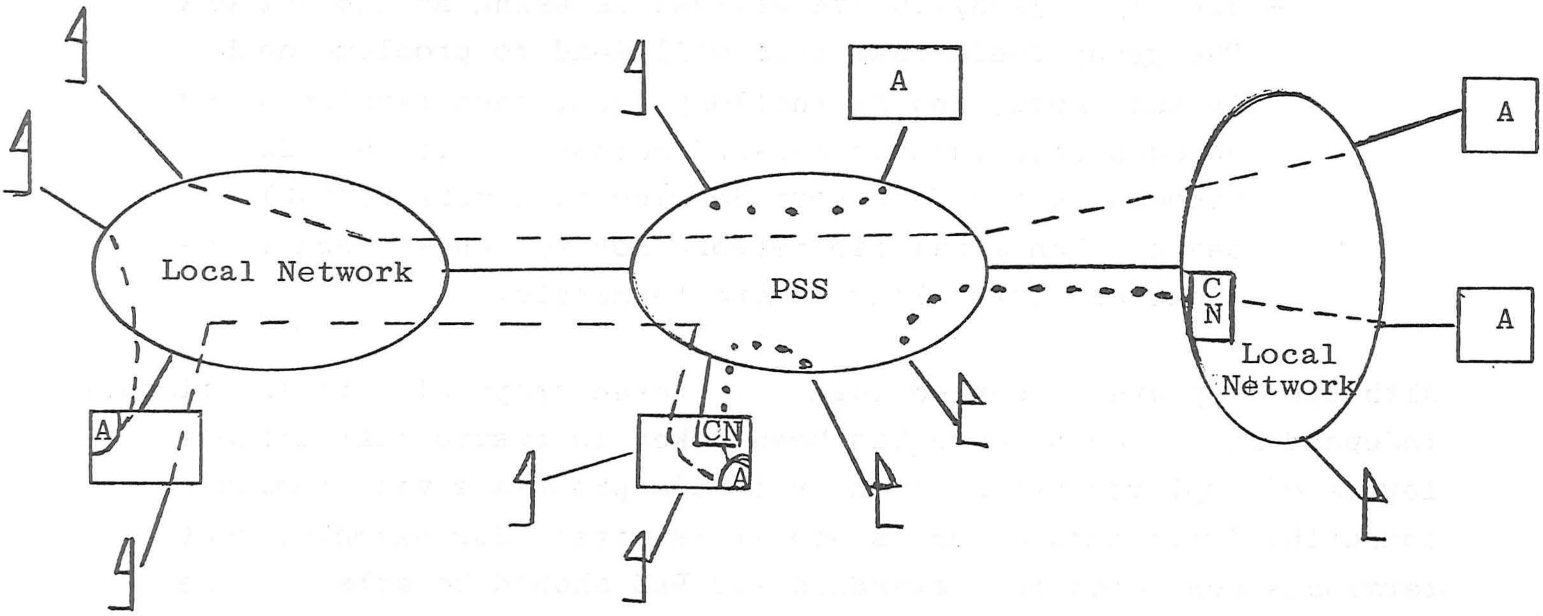

Terminals — Host computer

A Applications — CN Converter

---- TS29 connections — ...... X29sub connections

6. Future Progress

In the long term we can hope for standardisation of a Virtual Terminal Protocol by ISO or CCITT (or both) but in the short term there appears no justification to define one in competition with existing protocols such as the EURONET Data Entry Virtual Terminal Protocol. In the short term we must take into account the existence of X3/X28/X29 and the likely usage of the Transport Service over PSS.

In order to continue discussions with the Post Office on facilities and usages for PSS, a number of separate documents have been prepared and are attached to the report as Annexes. Annex A is a proposal for allocating TS29 parameters compatibly with existing usage, and Annex B is a detailed discussion of the proposed enhancements to X3/X28/X29 parameters and facilities; this we hope the Post Office will take into account in its enhancements to the PSS PAD and in discussions on future changes to the X3/X28/X29 Recommendations. Annex C proposes a set of profiles for PSS which would make X29sub available on the PSS PAD, as far as is possible, within our current understanding of the PSS implementation.

TERMINAL CONCENTRATORS, SWITCHES AND PADS

HARALD KIRKMAN

University of London Computer Centre

Terminal Concentrators, switches and PADs.

H. Kirkman,

University of London Computer Centre

A group of potential users and/or implementors of switching equipment for asynchronous terminals has been set up under the auspices of the JNT to examine the extent to which common requirements and standards for such devices can be agreed within the University/Research Council community. The existence of specifications for equipment of this sort, agreed by at least a subset of users, would be beneficial in that it would

provide input to manufacturers who may be contemplating the production of such equipment

ensure maximum compatability within 'home-made' implementations

permit the production (by a University or system house under contract) of a centrally-maintained 'kit of parts' - both hardware and software, from which specific configurations could be built.

The five classes of equipment so far identified and discussed by the group are as follows:

1. Contention unit.

This device has m 'input' ports for asynchronous terminals, and n (where n<m) 'output' ports to a host computer. Terminal users contend for the output ports on a 'first come - first served' basis.

2. Switching contention unit.

This device has m input ports and n sets of output ports to a number of different hosts. Typically the total number of output ports is less than the number of input ports. Users indicate the host they require (by typing an identifying character string for example) and then enter contention for the ports on that host as in (1) above.

3. Concentrator.

A concentrator has m input ports for asynchronous terminals and one (usually synchronous) output port for each of a number of hosts. Data from all the terminals connected to a given host must be multiplexed over the line to that host using a protocol that may be non-proprietary, but is usually manufacturer-specific. Concentrators (and PADs) have the advantage over the above devices in that additional terminals can be accomodated by the addition of input ports only, whereas the earlier devices potentially require further output ports, and input ports on the mainframes.

4. PAD.

A PAD (Packet Assembler/Disassembler) is similar to a concentrator, but is capable of connection to networks instead of, or as well as, host computers. As implied by its name it is assumed to be working in a packet-oriented environment. It is required to multiplex data over a single communications link for more that one host as well as for more that one terminal. In the case considered here, the protocol used on that link will be X.29.

5. Reverse PAD.

This device will accept multiplexed data (conforming to X.29) for a specific host and demultiplex it, feeding the data from different terminals into separate asynchronous ports on the host, thus simulating a number of locally-attached terminals. Such a device would allow existing systems to be connected to X.25 networks immediately with no need to modify the software or hardware on the host.

Devices of the first three types are, of course, already available from a number of manufacturers. Concentrators are not considered further, but there appears to be a definite need for a system which performs the switching and contention functions as an interim measure, capable of being upgraded later to a PAD (by changing the software and adding asynchronous port) as services based on triple-X become available.

Similarly, there is considerable interest in the concept of the reverse PAD as a means of providing early triple-X support on existing systems. The difficulty here lies in attempting to define and implement the diversity of protocols used by mainframes for conversing with asynchronous terminals.

There is, too, a clear need for private PADs, to act as concentrators using standard protocols, as switches providing access to a number of local hosts, and, in common with the PO PADs, providing access to a range of external services over X.25 networks.

Accordingly, the group is working on specfications for

1. a switching contention unit
2. a PAD
3. a reverse PAD.

It is intended that the same hardware components should be used for all the devices, and here two paths are open. One choice would be to produce special-purpose hardware based on LSI technology, which has the advantage of cheapness, but may prove difficult to maintain on a day-to-day basis. An alternative is to base the development on existing microcomputer architectures, but incurring a cost penalty as a result (say perhaps £400 per port as opposed to £100). The second course does offer the additional advantage that software developed for a standard system (e.g. LSI-11) would be available to the users of a large number of existing systems.

The specification of the switching contention unit is in preparation. The basic function is relatively easy to define, although some details, such as the mixing of full- and half-duplex hosts and connection/disconnection procedures require some care.

A specification for a PAD exists in draft form. The function of the device is, of course, broadly laid down in CCITT recommendations X.3/X.28/X.29, but reference is also made in the document to the Post Office implementation, and to the work of the Character Terminal Working Group of PO Study Group 3. The specification calls for a device capable of handling up to 32 asynchronous lines at up to 9.6 kbps (but typically with 8 terminals) with one or more synchronous connections at speeds up to 9.6 kbps.

Work has also started on the definition of the reverse PAD. On the network side this device is complementary to the PAD, and must therefore conform to X.29. As mentioned earlier, the host side is more difficult to specify in view of the wide variation in host-terminal protocols. It is hoped that definitions can be given for some of the widely-used hosts, and that a set of guidelines can be developed on how to tailor the device to map X.29 to other host interfaces.

When complete, these specifications will be distributed to industry, in the hope that some manufacturers may offer to supply such equipment (or may reconfigure existing equipment to fit). They will, of course, also be circulated within the University/Research Council community, with the intention (among others) of encouraging intending implementors of similar systems to consider bending their implementations slightly (if necessary), and offering them to the rest of the community through the JNT.

Anyone who would like to be kept up-to-date on the progress of this work, or who has input to contribute to the group is invited to get in touch with

H. Kirkman,
ULCC,
20 Guilford Street,
London
WC1N 1DZ

TERMINAL SWITCHES, CONCENTRATORS AND PADS.

GROUP SET UP BY JOINT NETWORK TEAM TO

"ELIMINATE UNNECESSARY DUPLICATION AND ENSURE WIDEST APPLICABILITY OF PRODUCTS"

BY

FORMULATING SPECIFICATIONS

DIVIDING UP WORK

1. CONTENTION UNIT

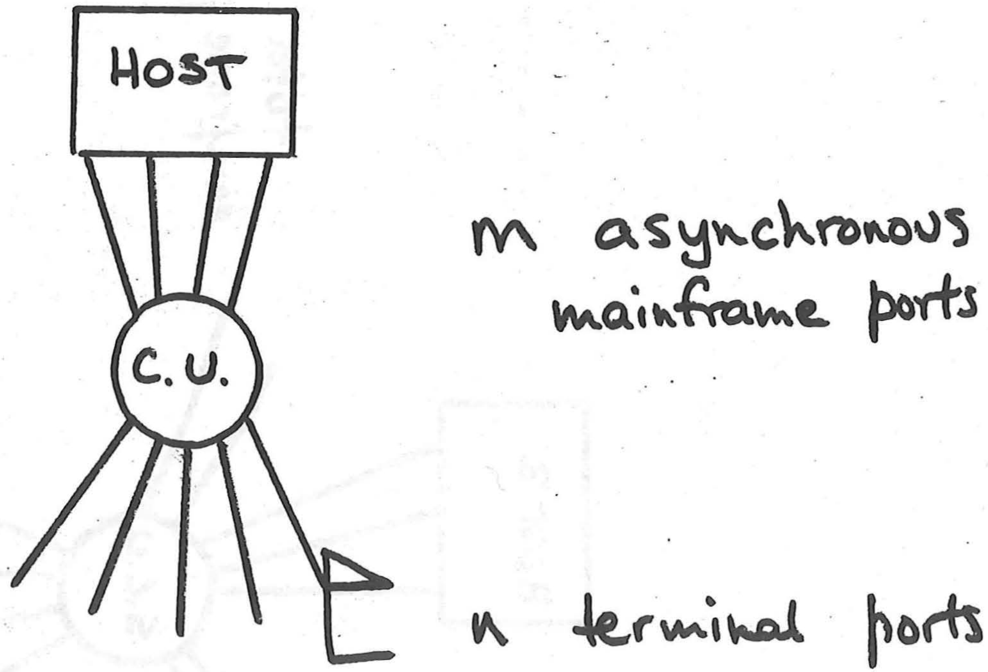

$m < n$

SWITCHING CONTENTION UNIT.

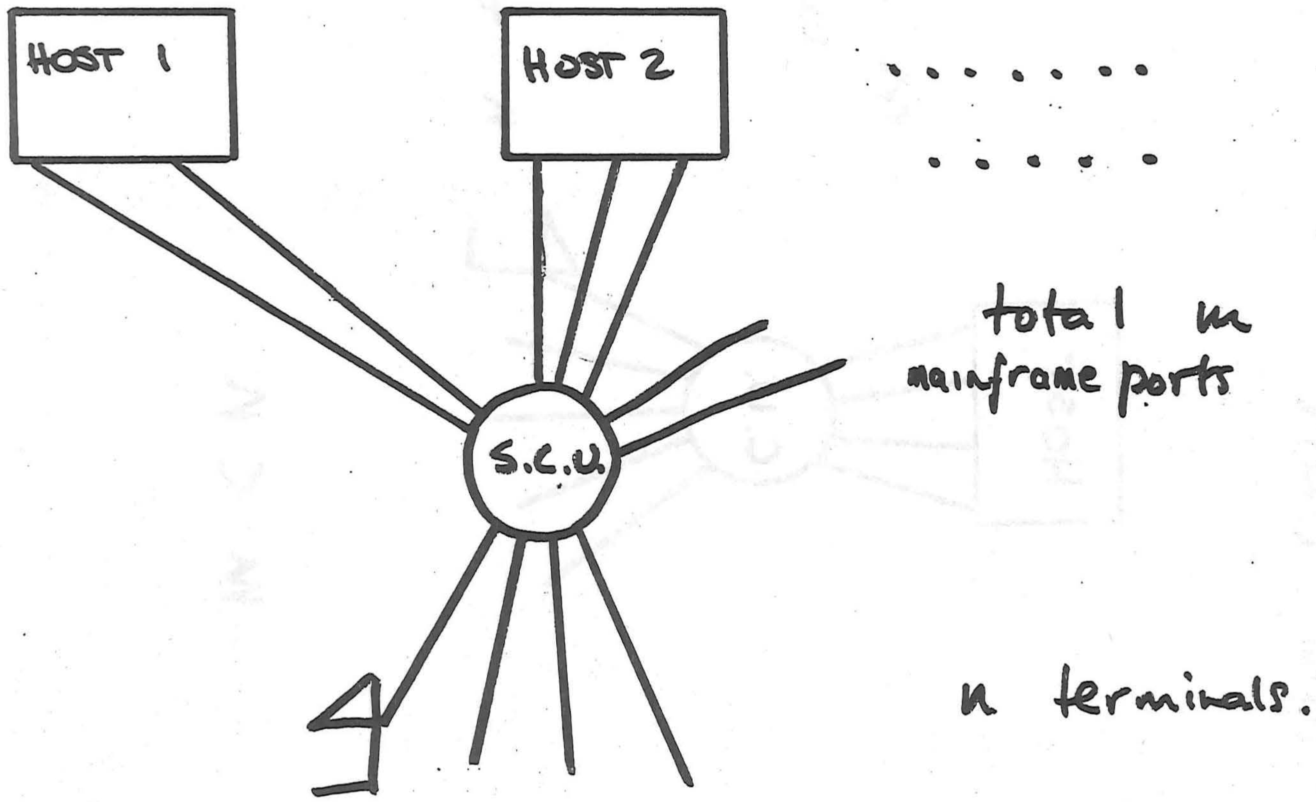

USER SELECTS HOST (BY TYPING CHARACTER STRING?) AND THEN ENTERS CONTENTION FOR THE PORTS ON THAT HOST.

(N.B. HOST MUST HAVE ONE PHYSICAL PORT FOR EACH ACTIVE USER).

3. CONCENTRATOR

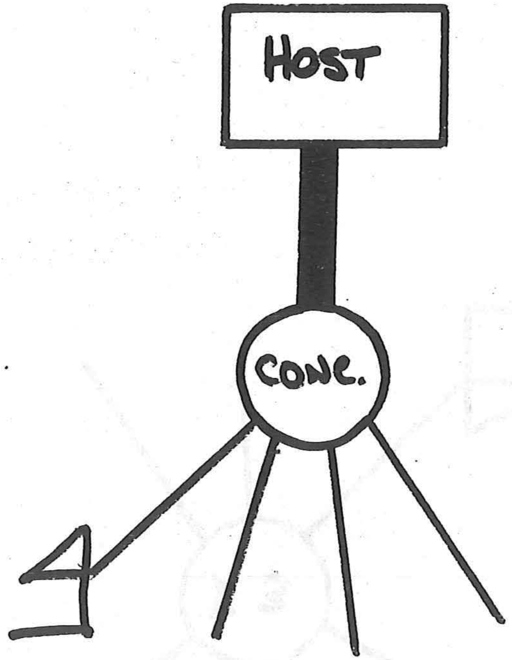

SYNCHRONOUS CONNECTION

n TERMINALS

ADDITIONAL USERS DO NOT REQUIRE ADDITIONAL MAINFRAME PORTS
(BUT PROPRIETARY PROTOCOLS OVER SYNCHRONOUS CONNECTION)

4. PAD. (PACKET ASSEMBLER/DISASSEMBLER)

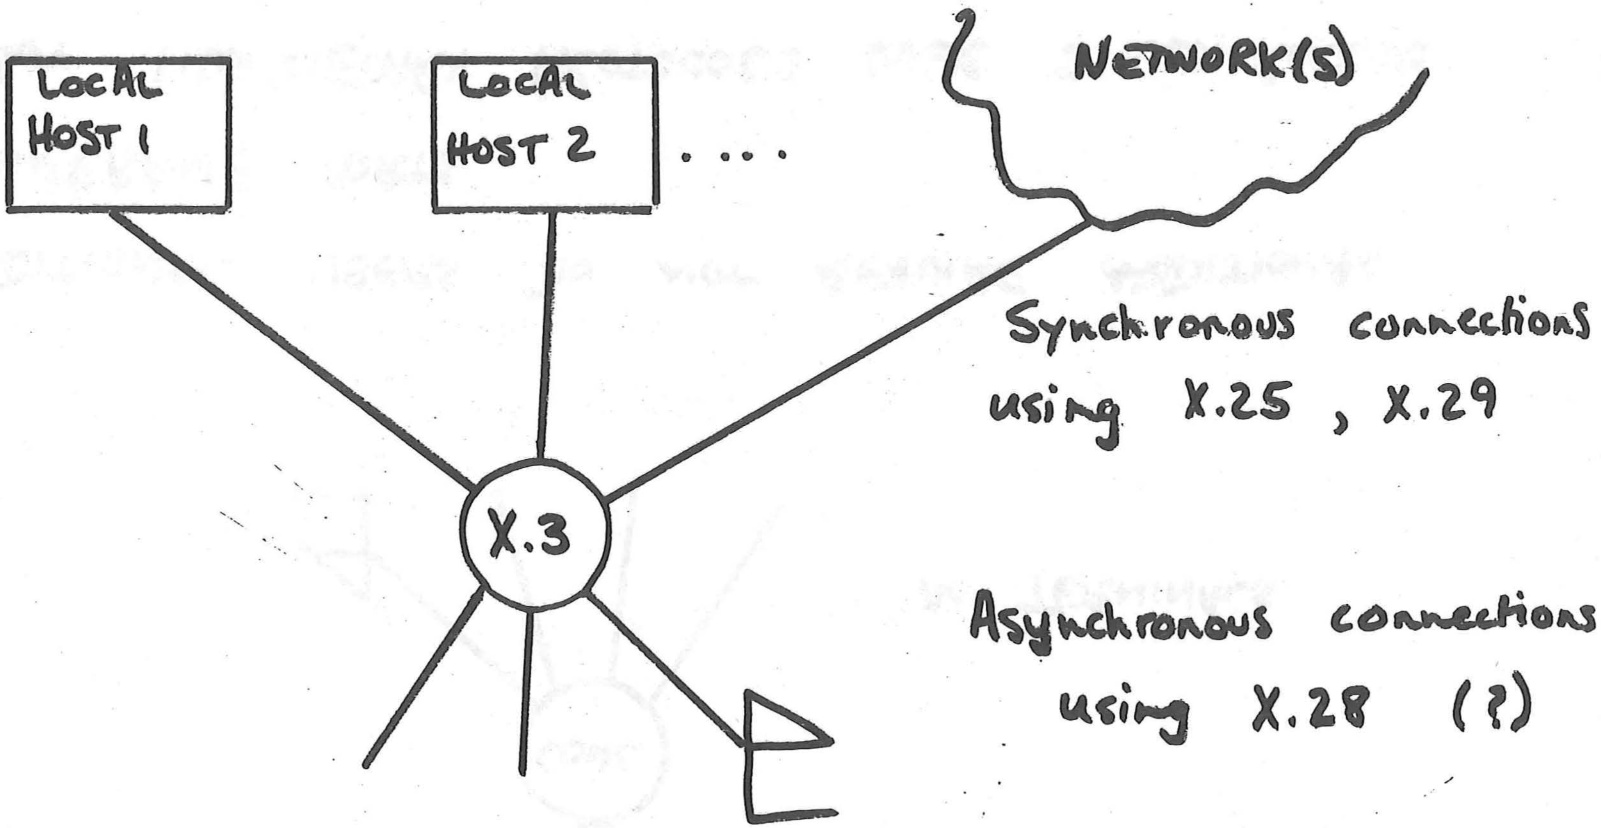

ᗡAꟼ

REVERSE PAD.

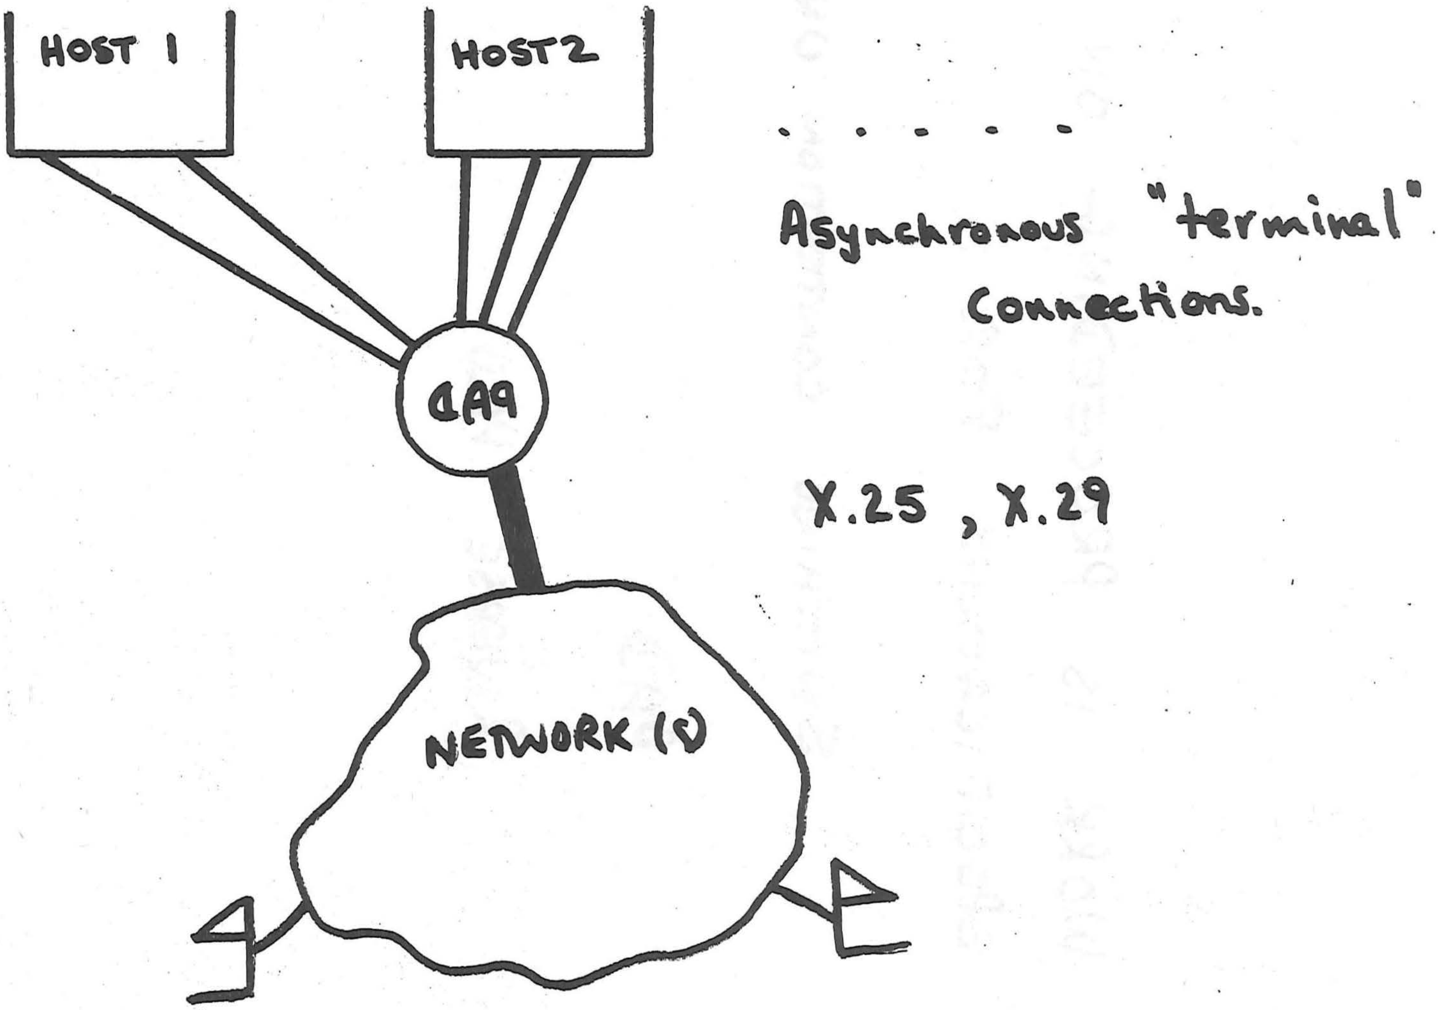

REVERSE PAD ALLOWS EXISTING HOSTS TO PROVIDE TRIPLE-X SERVICE WITH NO CHANGES TO HARD- OR SOFT-WARE.

WORK IS PROCEEDING ON SPECIFICATIONS FOR :

SWITCHING CONTENTION UNIT

PAD

REVERSE PAD.

GENERAL FEATURES.

ALL BASED ON SAME HARDWARE

Special purpose micro-based

Cheap

Difficult to arrange maintenance

General purpose μ-based (eg LSI-11)

Not cheap

Good maintenance

Programs useable in existing systems

SIZE (PAD, REVERSE PAD)

- UP TO 32 ASYNC CONNECTIONS AT UP TO 9.6 kbps, BUT TYPICALLY ~ 8.

- 1 OR 2 SYNCHRONOUS LINES AT UP TO 9.6 kbps

FUNCTION.

- SWITCHING / CONTENTION UNIT BEING DEFINED.

  (1/2 : FULL DUPLEX
  DISCONNECT SEQUENCE, ETC.)

- REVERSE PAD

  1/2 EASY : X:29

  1/2 DIFFICULT: HOST - SPECIFIC ASYNCHRONOUS PROTOCOL.

- PAD

  USING WORK OF CHARACTER TERMINAL WORKING GROUP OF PD STUDY GROUP 3.

PROGRESS ON THE RING AT KENT

MATT LEE and STEVE BINNS

University of Kent

REPORT OF THE HARDWARE PROGRESS ON THE UNIVERSITY OF KENT IMPLEMENTATION OF THE CAMBRIDGE RING

Hardware Situation at Networkshop 4

The Computing Laboratory started its implementation of the Cambridge Ring in July 1978, and below is a summary of the progress made between then and Networkshop 4.

1) (FEB. 1979) The minimum hardware necessary to support a RING packet structure was built, tested and run. This consisted of a Monitor Station, a Repeater, and a 100 metre drum of cable. With this arrangement, the Monitor Station was able to set up a valid RING structure, and maintain a check on it by sending randomly filled packets to itself.

2) Apart from this minimal RING, other hardware was being built, being sent away to be built, and being commissioned when it returned from being built. The RING REPEATERS were home-made, the RING WORK-STATIONS were assembled in the Laboratory Workshop and commercially wire-wrapped, and the RING PDP-11 INTERFACE LOGIC was being put on comercially bought general-purpose cards.

Progress from there ...

The Monitor Station, a Repeater, and the 100m drum of cable were all put on a trolley. This served as a test-bed for any hardware that needed commissioning. The ability to wheel the whole minimum RING around the laboratory was extremely useful, so much so that the an experimental service is now run with the Monitor Station still on the trolley. Modifications to the Monitor Station were made much simpler as the unit could be brought into the Workshop, and problems encountered while testing PDP-11 related hardware were similarly eased as the Monitor Station could be nearby.

With this portable test-bed, completed RING hardware was gradually tested. Below is a chronological sequence of the events to date.

1) (APRIL 1979) Networkshop 4 in York.

2) (MAY. 1979) The general-purpose interfaces used to connect the PDP-11s to the RING had a nasty error removed from the manufacturer's portion. Shortly

after the first PDP-11 interface was commissioned, and packets sent around the RING for the first time. The packets were sent from a PDP-11 to itself.

3) A simple box was designed to ease the physical connection of nodes to the RING, the physical connection of the RING cables themselves around the laboratory, and quick replacement of faulty hardware connected to the RING.

4) A 19" rack-mounting chassis was designed to house a RING NODE. The chassis allowed easy replacement of the enclosed REPEATER and WORK-STATION, and had a RING connection box fixed to it as well.

5) (JUNE 1979) A Motorola* M6800 Microprocessor interface was designed, built and tested, and was installed in the Laboratory's M6800 development system.

6) (JULY 1979) The specification for the Cambridge "Type-One" microprocessor system, including RING interface, was obtained from the Cambridge Computing Laboratory. The system is based on the Z80-type microprocessor.

7) (AUG. 1979) Other PDP-11 interfaces were commissioned, and programs written to test them.

8) University College London had developed some printed circuit boards for the REPEATERS and WORK-STATIONS, and some were bought for evaluation.

9) Cambridge Computing Laboratory supplied some details of an improved New MONITOR STATION, and work building this second unit was started.

10) Commissioned PDP-11 interfaces were quickly allocated to specific machines, and so some more were ordered.

The RING Connection System

One of the characteristics of typical computer laboratories is the state of the space where its many wires are hidden from view. If our RING was physically wired as a loop under the false floor we have in our laboratory, it would not be possible to remove it without breaking (which in the usual case would mean unscrewing or unsoldering) the RING. Doing this any number of times is undesirable, and so a simple connection box was devised which has proved to be versatile.

What are the desirable qualities of such a box?

1) Need to be able to add in nodes easily

2) Need to be able to remove nodes easily

3) Need to be able to extend cable lengths easily

4) Need to be able to find cable faults quickly

With these ideals in mind, the following box was designed. The unit consists of three D-type connectors mounted through the top of a small die-cast box and connected on the reverse side by a small printed circuit board. All of the connectors are alligned perpendicular to the longest side of the top of the box, and all 'point' in the same direction, from left to right, say. The three connectors are, from left to right, a 9-way plug, a 15-way plug and a 9-way socket. Signals from the RING are accepted by the 9-way plug, and are routed by the PCB to the appropriate pins of the 15-way plug in the centre of the box. From here, the signals go to a Repeater, say, and it sends refreshed and possibly new signals back down the cable to the remaining pins of the same 15-way plug. These returning signals are routed by the PCB onto the RING again via the 9-way socket.

Looking at some of the advantages of the box, I noticed

1) Inserting a node into the RING was easy. A nearby RING cable is cut, a 9-way socket put on the active end, and a 9-way plug on the passive end. Notice that these two can be plugged into one another and the RING will again work. These two 9-way connectors are then plugged into the appropriate sections on the top of the box. A 4-pair cable with a 15-way socket <u>on both ends</u> then connects the centre plug of the box to the Repeater of the node that is being connected.

2) Removing a node from the RING involves unplugging the two 9-way connectors from the box and joining them together. The previously connected node is now isolated.

3) Only two different types of cable need ever be made up in the workshop. One is the one-way RING cable, and the other the two-way RING cable.

4) The cables made up for joining boxes together may be used, and re-used anywhere in the RING. This is useful, and economic, if equipment needs to be moved round the laboratory and re-connected to the RING. The one-way cables can, of course, be plugged into

one-another to make longer extensions.

5) Re-routing cables around the laboratory is made simple, as each end of every cable may be easily freed.

Another interesting point was then noticed. The boxes themselves could be connected together by the centre 15-way plugs. This meant that, using a length of 4-pair cable, a RING SPUR, say, could be run to another room.

Once this was realised, the possibilities were endless, and bounded by the sole condition that it is currently standard to have no more than 100m of cable between any two Repeaters.

A number of small diagrams over the page describe some of the more useful configurations of the RING using the box. It has proved very convenient in this project to have such a flexible connection system, especially during the initial stages.

University College London RING Hardware

The original specification for the Kent RING was to copy the Cambridge hardware exactly. Now, however, Steve Wilbur at University College London, has produced the two essential parts of any RING node on printed circuit boards, and a few sets of these PCBs have been bought and are being tried in the Kent RING.

The two sets available from Steve Wilbur are:

1) A Repeater. This consists of only one board, which contains the interconnections for the repeater power supply and modulation/demodulation logic.

2) A Work-Station. This consists of three boards, sensibly split into a Timing and Data Serialisation Board, a Transmit Logic Board, and a Receive Logic Board.

The Timing and Data Serialisation Board is useful as it provides the necessary timing signals for any of the likely interesting points that occur in a passing packet. It forms a good basis for any specific application needing these timing waveforms.

Each of the boards is double sided, but at present not plated-through-hole. Should sufficient interest be taken, a new set of plated-through boards will be produced for both the Repeater and Work-Station. Each board currently costs around twenty-five pounds, and therefore a complete set of three boards for a Work-Station will cost around seventy-five pounds. Interested parties should contact

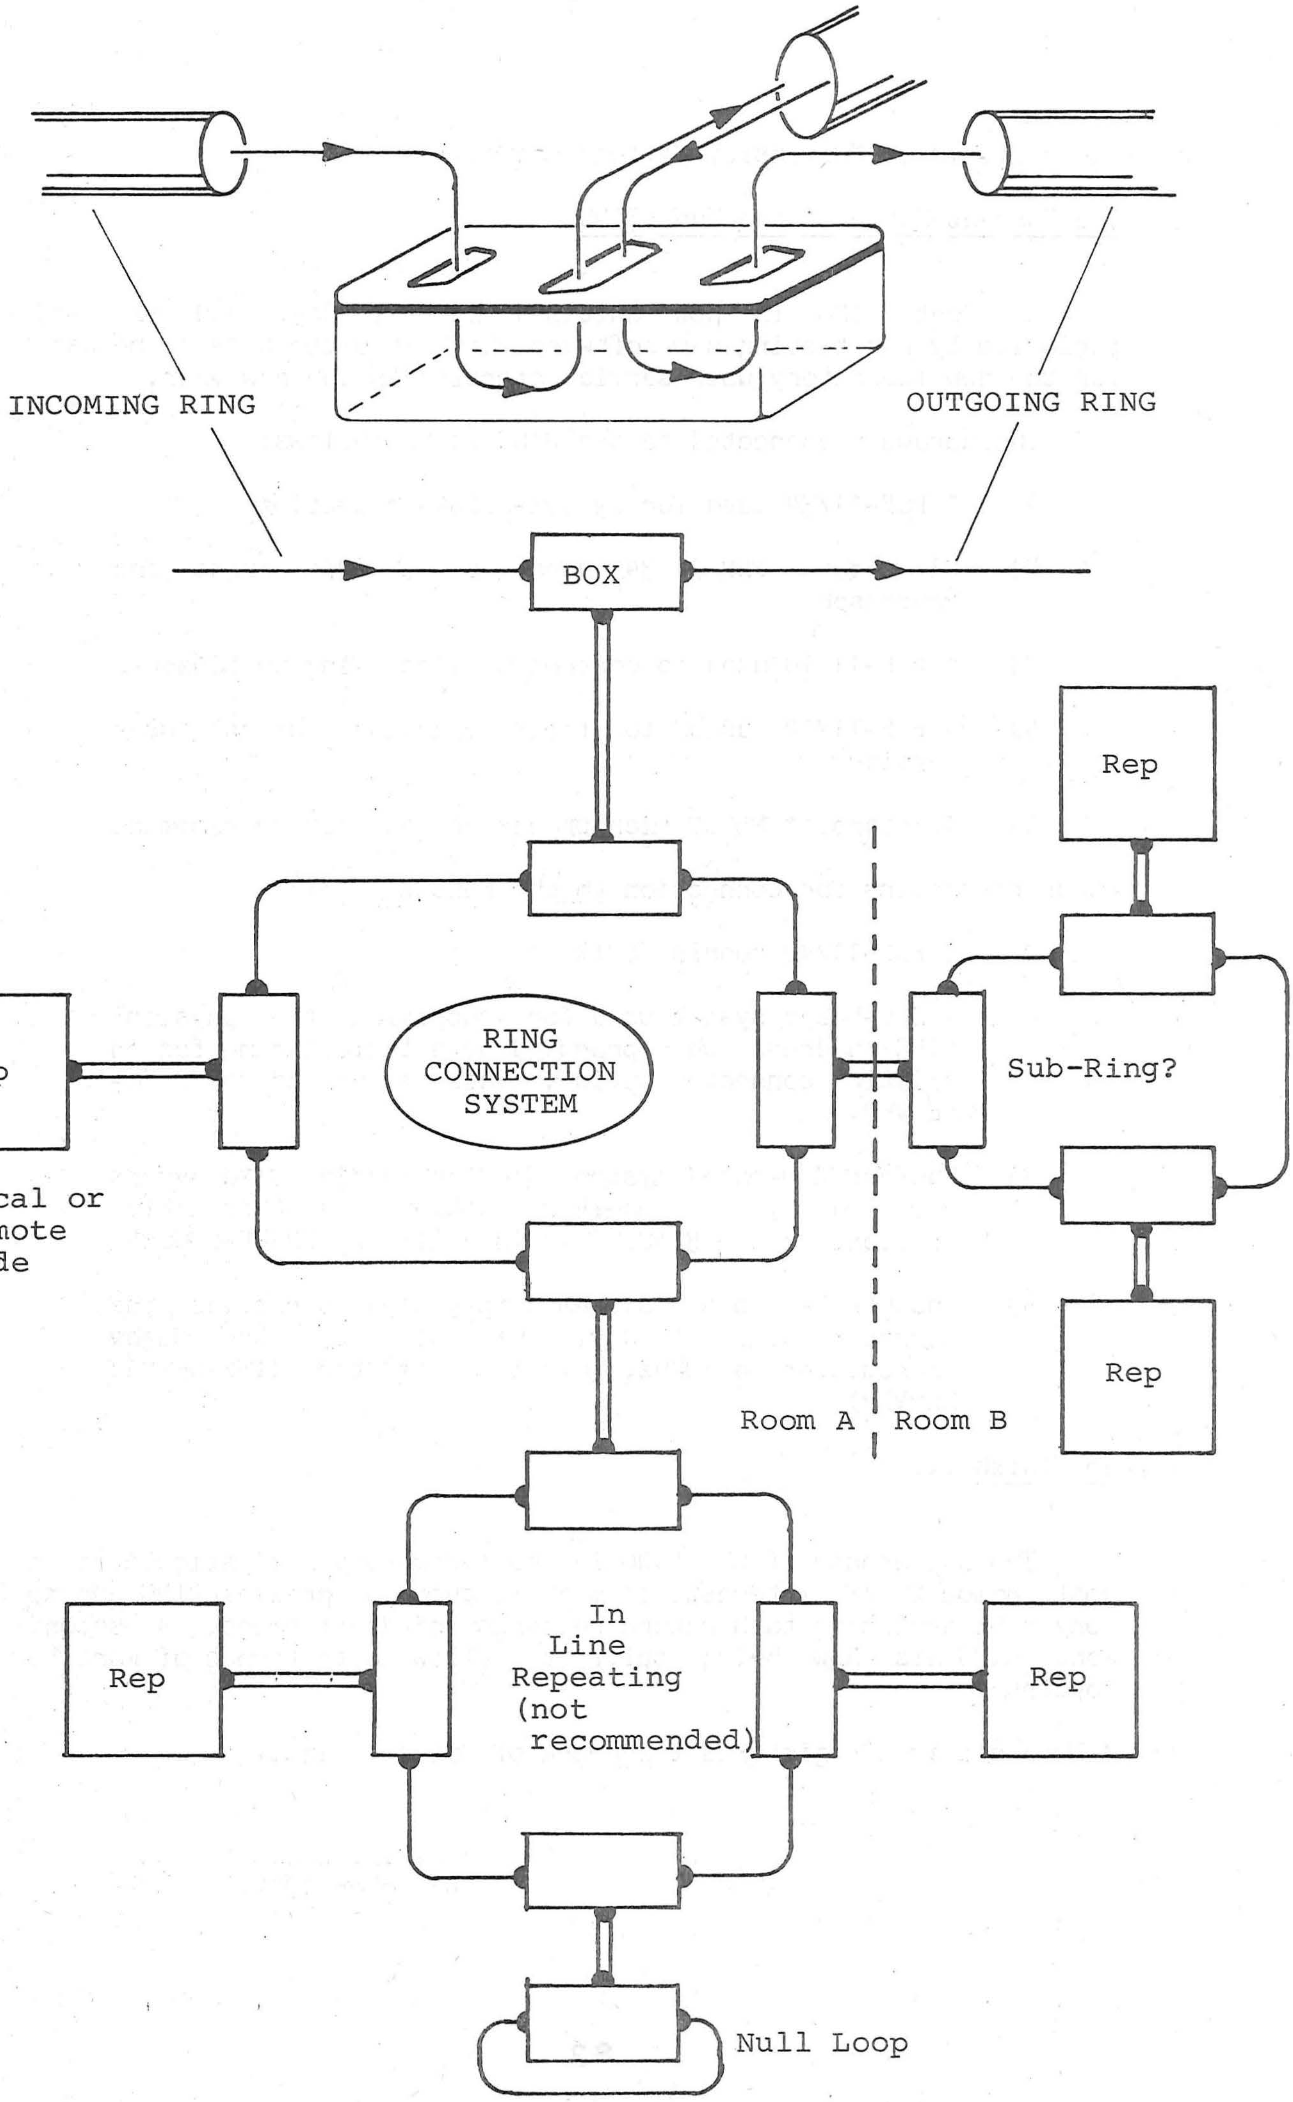
INCOMING RING
OUTGOING RING
BOX
Rep
RING
CONNECTION
SYSTEM
Local or
Remote
Node
Rep
Sub-Ring?
Rep
Room A
Room B
Rep
In
Line
Repeating
(not
recommended)
Rep
Null Loop

Steve Wilbur at University College London.

The Current State of the Kent RING

The Kent RING is now switched on all day, and is used periodically for testing new software which is ultimately to be used for the new laboratory user service starting in the new year.

The hardware connected to the RING is as follows:

1) A PDP-11/34 used for system-software testing

2) A second PDP-11/34 used as an EMAS Front End Processor

3) A PDP-11/10 used to control a direct link to London

4) A PDP-11/10 used to front end the current user service

5) A Motorola* M6800 microprocessor used for development

Other contenders for connection to the Kent RING are:

1) A PDP-11/40 running UNIX

2) A Z80-based system used for supplying the physical RING address when prompted by a logical name for an existing connected device. This is called a NAME-SERVER.

3) Another Z80-based system used for logging RING errors detected by the Monitor Station, and other active Stations on the RING. This is called an ERROR-LOGGER.

4) Another Z-80 based system simply used to provide the correct time of day. This will use the Rugby transmitter on 60KHz. This is called the TIME-OF-DAY SERVER.

To Finish ...

The appearance of the RING in the laboratory has stimulated a vast ammount of interest, so much so that the present RING can no longer be dedicated to hardware research and development. A second Kent RING is now being built to allow both areas of work to continue.

* Motorola is a registered trademark of Motorola Inc..

M.N.A. Lee,
November 1979.

Putting the Ring into Service - Software

The main computing service machine at Kent is currently an ICL 2960 running VME/K. At the start of next year we hope to install a new interactive machine to enhance the multi-access facilities. Currently a DEC 11/40 and a DEC 11/34 provide a UNIX service for computing science teaching and research. Also the final elements of a prototype Cambridge ring system are being commissioned within the Computing Laboratory.

Earlier this year a change of operating system on the 2960 and the future use of the ring were separately under discussion. For the ring the likely possibility was a gradual introduction of facilities, terminal concentrators for example and the implementation of job and file transfer protocols. It was strongly felt that the Computing Laboratory must give a positive lead in the implementation of high level protocols before the ring could be made available for general use throughout the University. Although the ring has the advantage that two stations may communicate with a stream of raw data or some non-standard protocol without effecting other users of the ring system, the use of pertinent protocols is essential if the ring is to be used as a mechanism for sharing resources, particularly if gateways to other local ring systems and external networks are provided.

One possibility for a new operating system for the 2960 was EMAS as developed at the University of Edinburgh. This system seemed to offer the facilities and features required for a University service environment. However, at Edinburgh the communications are handled through RCONET a network of PDP11 based nodes. RCONET did not seem relevant to Kent and there were reservations on the consequential level of funding required.

As is perhaps obvious in retrospect, we decided to solve both problems by using the ring as the basis of a communications network for an EMAS system on the 2960. Currently we hope to offer a full EMAS service after Christmas 1979. The main devices on the ring will then be as in the diagram below.

Because of the short timescale involved the bulk of the EMAS software must remain unchanged, this will include the RCONET interactive terminal and RJE protocols. Beneath these protocols there is a sophisticated transport service known as NSI (node standard interface). To minimise software changes this is embedded in the ring byte stream protocol (BSP) which is an error correcting transmission protocol designed at Cambridge. The required software changes are to the EMAS front end, the Edinburgh terminal concentrator, together with substantial modifications to the Kent RJE emulator used for ULCC and the 2960 system. A byte stream protocol exerciser has been written to aid the debugging of these items. A highly desirable piece of software to test the whole system is the Edinburgh terminal emulator (ERTE). This has been modified to use the ring and we hope to be able to simulate a heavy terminal load on the 2960 before the full service is mounted.

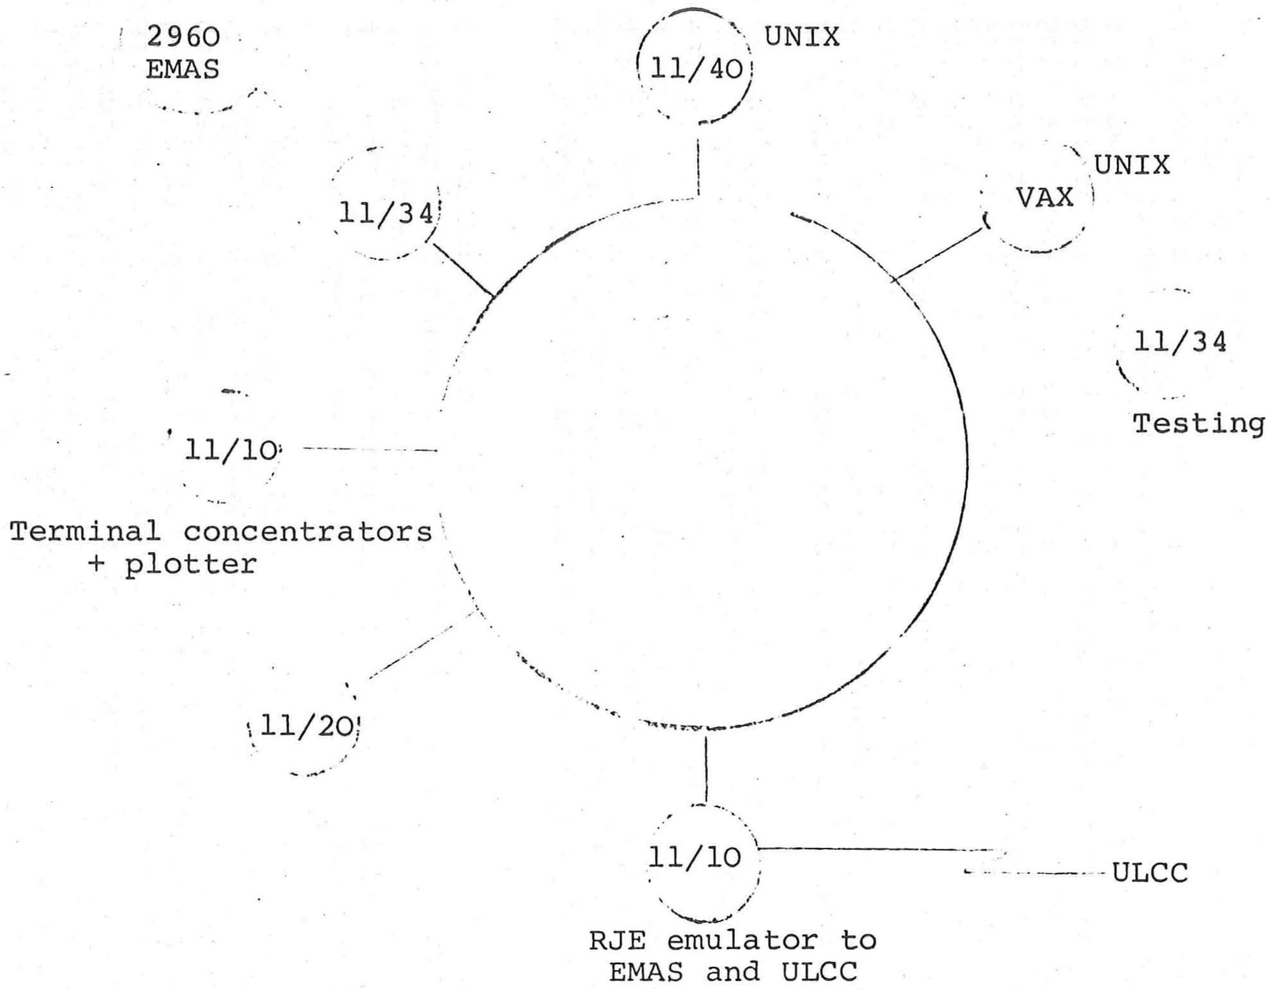
2960
EMAS
11/40
UNIX
VAX
UNIX
11/34
11/34
Testing
11/10
Terminal concentrators
+ plotter
11/20
11/10
ULCC
RJE emulator to
EMAS and ULCC

One obvious disadvantage of this course of action is that the extension of the ring and use of the ring as a research tool might be adversely affected by the heavy dependence of the computing service on the ring. However this is balanced by the fact that computing laboratory resources are concentrated on the ring and not split between service facilities and research facilities. In particular the provision of a common transfer mechanism, using BSP, between all laboratory machines is a considerable gain. A simple file transfer mechanism is being implemented for the UNIX system, and the ring is already being used for down-line loading of various front end systems.

We think it is important that the use of the ring for a full computing service should be demonstrated with high level protocols. The replacement of the RCONET protocols by the de facto standard protocols that will emerge, based on the BSP as a transport service is clearly an important priority. However this would not seriously effect the service to users via the ring as the actual protocols used are user transparent. A first step will probably be to replace NSI by the byte stream protocol.

The integration of the new interactive machine (probably running UNIX) into the ring system will probably be done using the new interactive protocol. One terminal feature that UNIX users will demand is a 'raw mode' input option selectable by the host as this is used for some UNIX programs.

Extension of the ring to the whole campus will depend on the provision of ring bridges, we can then concentrate on providing shareable facilities and resources and ring expertise for the whole user community.

18.10.79

TRENDS IN COMMUNICATIONS HARDWARE

RAY CHISHOLM

Edinburgh Regional Computer Centre

TRENDS IN COMMUNICATIONS HARDWARE

There have been a number of developments in hardware based technology which will impact significantly on a data communications scene which is now maturing very rapidly as a result of international agreement on networking standards. This report attempts to highlight some of the developments which I consider most relevant to our activities.

NEW CCITT ELECTRICAL INTERFACE STANDARDS

The ubiquitous V24 interface, which complies with electrical interface specification V28, is in this age of micros, fibre-optic and digital data network technology beginning to show its age - the 20 Kbits/sec maximum data rate goes nowhere near supporting the demands of these technologies. Even the V35 interface standard which is specified to cover data rates up to 60 Kbits/sec falls well short of the perceived requirements. Two new standards have been defined which provide a reasonable transition path from the modest requirements of yesterday to the megabit speeds of tomorrow.

X26

* Unbalanced circuit operation.
* Capable of inter-operation with V28 and X27 standards (Bridges old-to-new).
* Supports transmission at 300K Kbits/sec over 40 feet (c.f. V28-20 Kbits/sec over 50 feet).
* Also known as V10; some as EIA RS423.
* Based around I/C technology.

X27

* Balanced circuit operation.
* Supports transmission at 10M bits/sec over 40 feet - 100 Kbits/sec over 3000 feet.
* Also known as V11; same as EIA RS422.
* Based around I/C technology.

It should be noted that many of the micro based systems on the market are offering interfaces to these electrical standards. The EIA (Americans) have defined a functional and mechanical interface RS449 which incorporates the RS422/423 electrical characteristics.

OPTICAL FIBRE DATA TRANSMISSION SYSTEMS

This last year has seen the introduction of a number of fibre optic transmission systems which can be used to replace conventional modem based transmission links. Many/

Many of these systems are offered with RS232C (V24) or RS422 type interfaces. They are configured as stand-alone or rack mounted transmission equipment interconnected by a fibre-optic cable.

FIBRE OPTICS OFFER -

* Bandwidth - distance product $\underset{\sim}{\sim}$ 100 Gigabit - metre/sec.
* Low Bit Error Rate (BER) $\underset{\sim}{\sim} 10^{-9}$.
* Capable of supporting data rates of 10 Mbits/sec over a 2 km link.
* Smaller conductor.
* Elimination of ground loops.
* No radiated interference.
* Immunity to electro-magnetic and radio interference.
* Effective bandwidth orders of magnitude greater then copper.

The Post Office are currently experimenting with fibre optic systems and have installed repeater - less links running at 8 M bits/sec over 12 km and 140M bits/sec over 8 km. Another interesting fibre product development is a device known as a transmissive star coupler which will find applications in the local area networking arena. It is a passive device to which up to 32 ports (16TX & 16RX) can be coupled, each connected station has a transmit and receive port connection and any transmit port distributes its optical signal evenly to all receive ports. A local area network based on this type of technology is being developed at Xerox, Palo Alto - this system is known as FIBERNET.

'BLACK-BOX' INTERFACES

A previous speaker posed the question "How does one connect old communications equipment using old protocols to new networks using new protocols?". The answer, courtesy of micro based technology, is to use a black box interface to perform the functional translations necessary. There are two approaches being adopted, one where the black box is external to the terminal equipment and connects to it using a conventional data communications port connection, the other is where the black box is parasitic and resides within the terminal system connecting directly to the internal bus of the processor.

An example of the first type of device is the MPAC 5000.

MPAC 5000

* Z80 microcomputer based.
* 2 - X25 Network ports and 4-16 terminal ports.
* Firmware/

* Firmware supports X25 levels 1-3 and 'transport level' support available.
* Data compression & encryption options available.
* Supports a variety of user terminal types.

This product is manufactured by a Canadian company Memotec Services Corp. and available through Nolton Communications in the U.K.

An example of the second type of device is the UMC-Z80 which is Z80 microprocessor system contained on a PDP11 Unibus hex width module.

UMC-Z80

* Processor board also includes 20K bytes memory, 2 full duplex Async/sync data channels and UNIBUS DMA logic.
* Memory expansion board provides up to 64 K bytes RAM and 32 K bytes ROM.
* Serial line expansion board supports up to 16 duplex.lines. Each pair of lines has a dedicated Z80 CPU with local memory as well as access to global memory.

This product is manufactured by Associated Computer Consultants. Santa Barbara, USA.

X25 PACKET NETWORK INTERFACE I/C

An American integrated circuit manufacturer (Western Digital) is currently developing an LSI chip which will perform X25 packet network control to level 1 & 2 specification, LAPA or LAPB.

WD2501

* Serial line speed to 1.6M bits/sec.
* 48 Pin package.
* Programmable primary timer, retransmission counter and A-field.
* Uses DMA for transmission and reception of data.
* Contains 3 micro controllers.
* Zero bit insertion and deletion.
* Automatic FCS appendage and checking

It is my understanding that this device is currently sampling in the USA with initial batch qualities scheduled for the first quarter of next year. It's my guess that those who wish to be first in will be paying £250 - £300/chip.

CONTENTION BUS LOCAL AREA NETWORKS/

CONTENTION BUS LOCAL AREA NETWORKS

Ethernet is the best known example of a contention bus local area network but there are other systems becoming commercially available, particularly in the USA, which are variants of this basic technology.

NETWORK SYSTEMS CORPORATION

* Product known as HYPERCHANNEL.
* Coaxial cable based. Transmission rate from 50M bit/sec at 1000 feet to 1.5M bits/sec at 5000 feet.
* Adaptors available for a range of mainframes minis, and peripherals.
* Each adaptor can have up to four separate network connections.
* Prices! £25,000 / adaptor.

DIGITAL COMMUNICATIONS CORPORATION

* Product known as Bus Interface Unit.
* Coaxial cable based transmission at 1 Mbit/sec.
* Black box unit with a network connection and 2 user interface connections each capable of serial operation in ASYNC/SYNC mode up to 56 Kbits/sec.
* Uses HDLC on bus.
* Z80 based with 16K Rom and 8K Rom.

THREE RIVERS COMPUTER CORPORATION

* Computer system which is network based.
* 10M bit/sec coxial cable up to 2000 feet with 64 network ports.
* Interface to PDP11 being developed.
* Product known as PERQ network.

There are a number of other products which are coxial cable based which use a polled bus technique instead of the distributed control that these systems employ.

It should be fairly obvious from this rather cursory glance at a few of the developments on the hardware front that technology is advancing at a prodigious rate and still showing no signs of the predicted slow down.

R. CHISHOLM
EDINBURGH REGIONAL COMPUTING CENTRE

OCTOBER, 1979

SURVEY OF DEC COMMUNICATION PRODUCTS

IAN SERVICE

University of York

U N I V E R S I T Y   O F   Y O R K

DEPARTMENT OF COMPUTER SCIENCE

Survey of DEC communication products

This is a rather vague title (probably for a rather vague talk) and I have chosen to interpret it as applying only the DEC PDP-11's as all the good DEC computers (i.e. DEC-10's) use PDP-11's as communication processors.

In fact, the PDP-11 is used in many systems as the communications processor and its communications interfaces deserve some discussion. What I have chosen to do is to outline a few interfaces, books, and pieces of software I know about and I would ask anybody who knows of any I have forgotten to let me know, and I will either include a note on them in the proceedings or in a report to the next meeting of the PDP-11 networks group. (And if you do not know about that group contact Paul Kummer or Roland).

I have not concentrated on DEC-only equipment, though the majority is. I have no objection to PCM if it is cheaper or better, but there does not seem to be much in the synchronous interface line - perhaps the market is not big enough.

First of all, the books that you can refer to to find out about DEC communications products:

This slide shows the books I use. Most of these do not have any code number on them so you can only obtain them by berating your DEC salesman.

Terminals and Communications Handbook - 1979

- Essential to anybody doing PDP-11 communications but NOT reliable on detail and definitely not complete.

PDP-11 Peripherals Handbook

- Getting a bit out of date now that DEC appear to have split this into several books.

Distributed Systems Handbook

- Not very interesting

DEC PDP-11 books

Fig. 1/Slide 1

Now the interfaces you can use for synchronous communications on a PDP-11:

DU-11	- Obsolete (?)
DUV-11	- LSI-11 only
DUP-11	- Non DMA/hardware HDLC framing
DQ-11	- DMA
DV-11	- Multiplexed lines (up to 16) Asynch. or Synch.
DQS-11*	- Hardware HDLC/Dual interface
DMS-11F*	- LAP Framing in microcode

* - DEC CSS products

<u>DEC PDP-11 synchronous interfaces</u>

Fig. 2/Slide 2

Thus the interfaces which are most suitable for using are:

DUP-11

*DQS-11

*DMC-11F

Fig. 3/Slide 3

and as the last two are CSS products they tend to be rather expensive.

The DUP-11 is therefore the favourite standard interface for HDLC work on PDP-11's, but its big disadvantage is that it is an interrupt by character device and though it performs flag and CRC generation and checking by hardware, it still requires a lot of processor power to drive it.

It is possible to turn a DUP-11 into a DMA-like device by using a KMC-11. This is a microprocessor that sits on the PDP-11 unibus and polls the DUP-11 along the UNIBUS. This device raises PDP-11 interrupts only when the DUP-11 has finished prosessing a message, rather than when it has finished processing a character. This device frees up PDP-11 processor power, though I understand it can generate a lot of UNIBUS traffic. Standard software is available for the KMC-11 in a package called COM-IOP DUP which provides KMC-11 microcode for driving up to 16 (?) DUP-11's.

It is worth noting that the DMS11-F is in fact a KMC-11 with special synchronous interfaces on the microbus of the KMC-11 microprocessor, and of course some new microcode. This seems a more desirable solution, but as I said earlier it is made by DEC CSS (France) and supports only LAP under RSX at the moment. It also is rather expensive for a single line version but may be configured to support up to eight lines. Obviously one could reprogram the

the microcode (a development system is available) but the size of the RAM on the KMC-11A in only 1K, so there is not much space.

That's all the DEC synchronous equipment that I know about, and there are not really very many suitable for bit oriented protocols such as HDLC.

If we move onto non DEC interfaces for PDP-11's there are not many more:

UCL interface (no details - may be being marketed?)
UMC-z8Ø - by A.C.T.

Fig. 4/Slide 4

The UMC interface looks a very sexy device with a high degree of flexibility. If anybody knows of any others I would be glad to hear of them.

That's all I want to say for the moment and I apologise to all the other DEC systems that I have neglected.

J D Service
26 September 1979

Survey of DEC communications products - Supplement

Since giving my talk at Networkshop I have been informed about two other PDP-11 synchronous interfaces:-

1 This is a DUP-11 like interface which UCL have made for Q bus PDP-11's (i.e. LSI-11's). For further details contact:-

Peter L Higginson
Department of Statistics and Computer Science
University College London
Gower Street
London

2 A DMA interface for the PDP-11 unibus which uses a Z80 and a 5025 chip. For further details contact:-

Ahmed Patel
or Michael Purser
or Chris Horn
or Bryan Alton
Computer Science Department
Trinity College
201 Pearse Street
Dublin 2
Eire
Tel: DUBLIN 772941 Ext 1765

THE FILE TRANSFER PROTOCOL - PROBLEMS AND PROGRESS

DAVE RAYNER

National Physical Laboratory

File Transfer Protocol - Progress and Problems

by D. Rayner, NPL
20 September 1979

Introduction

The Network Independent File Transfer Protocol (FTP), as specified by the "Blue Book" [1], has been around since December 1977. This paper reports on the progress made in implementing and testing the protocol at NPL as part of the cooperative testing that has been carried out using EPSS. This paper also reports on the work of the FTP Implementors' Group in maintaining the protocol as an interim standard and the work of BSI DPS 20 (Open Systems Interconnection) Working Group 6 (WG6) on the development of a proposal for a longer-term standard to be submitted to ISO TC97 SC16 (Open Systems Interconnection). In doing this the major problems with the current FTP will be identified.

An Outline of the Protocol

FTP has commands at two levels. The first, Level 0, concerns the negotiation of attributes and termination. At this level the two communicating processes are known as P, the initiator, and Q, the responder. P has three commands:

SFT - to start a file transfer and specify its requirements for the attributes;
GO - to enter the data transfer phase;
and STOP - to terminate the exchange of commands.

Q has two commands:

RPOS - to give a positive reply to the SFT and specify any changes it would like to see in the attribute values;
and RNEG - to give a negative reply to the SFT and specify why it is rejecting the transfer attempt.

When P receives an RPOS it has the choice of accepting it with a GO or rejecting it with a STOP.

There are 25 attributes, all of which may be specified by parameters in the SFT. They are composed of 13 transfer control attributes, which specify details of the actual transfer such as the maximum record size, 7 identification attributes, which are used to select the desired file and specify information for checking access rights, and 5 miscellaneous ones, including 2 for carrying text messages. It is only really the transfer control attributes that are subject to negotiation. An important one of these is the Facilities attribute which allows for the specification of 6 optional facilities: data compression, retransmission of part of the file in a new transfer, restarts from the last acknowledged mark point in the same transfer, mandatory mark acknowledgement, temporary pause in the data flow at receiver request, and parameters on GO.

The second level, Level 1, concerns the data transfer phase. At this level the two communicating processes are S, the sender, and R, the receiver. Each has four available commands. Those of S are:

SS - to mark the start of the data

CS - to select the code to be used
MS - to indicate a mark in the data
and ES - to indicate the temporary or permanent end of the data.

Those of R are:

RR - to request a restart
QR - to request the temporary or permanent end of the data flow
MR - to acknowledge the receipt of a mark
and ER - to acknowledge the end of the data.

Implementation Progress

There are now 5 implementations of FTP active in cooperative testing over EPSS. These are at CADC, NPL, SWURCC, UCL and AERE Harwell, given in the order in which NPL became aware of their readiness for inter-site testing.

At NPL, processes P, S and R were all written by July 1978, but testing did not start until December 1978 because of a delay in getting the EPSS Transport Service, called Bridge, working. The first successful transfer occurred in January 1979 and by March transfers to and from CADC had become reliable and were being used for real traffic. The implementation of process Q was delayed by the inter-site testing of P, so back-to-back testing did not start until September 1979.

Implementation Size and Effort Required

It is difficult to give an accurate guide as to the likely size of an FTP implementation or the amount of effort required to produce one. This is partly because the NPL implementation contains various monitoring aids which would not be present in a production version, but more significant is the fact that the operating system environment and the complexity required in the negotiation can vary greatly and have a major effect on the size and timescale. Nevertheless, it is possible to estimate the likely ranges. The size is likely to be between 6 and 16 K instructions plus 4 to 24 K bytes of data space. The effort required is likely to be between 2 and 24 man months.

Maintenance of the Interim Standard

The testing of FTP on EPSS is a necessary preparation for its use as an interim standard, in particular for use on PSS. An interim standard is needed because of the slow progress towards an internationally agreed long-term standard. Currently this interim standard is being maintained by the DCPU FTP Implementors' Group, which has produced a list of recommendations and reminders to implementors [2] to clarify the specification and resolve ambiguities where these are found. This list is updated quarterly and distributed to all those on the Group's mailing list. In addition, the latest version will be enclosed with any future mailings of the "Blue Book".

The Group has also produced a list of suggested extensions [3] to overcome deficiencies in the protocol. The intention is to maintain the stability of the interim standard, so that all the suggested extensions are upwards compatible and will require acceptance by a wider community before any implementors are recommended to adopt them. Perhaps their main use at present is as input to the work of BSI DPS 20 WG6 towards a long-term standard.

Recommendations

18 recommendations have been produced to date, 9 agreed and 9 in draft. The breakdown into topics is as follows:

3 on qualifiers on RPOS, RNEG, STOP and GO;
4 on parameters on RNEG, STOP and GO;
2 on the use of operators;
3 on format effectors;
2 on codes;
3 on binary mapping;
and 1 on messages.

It is significant that all of these concern areas of the protocol that are specified only in English and most of them concern the negotiation phase.

The specification concentrates on those parts of the protocol that are needed for successful transfers, so it is not surprising that major recommendations have had to be made to tighten up the protocol in cases of rejection. It may be felt that such areas are of little importance, but on the contrary it is very important to the user who gets his transfer request rejected to know why it has been rejected. Unfortunately, ambiguities in the specification can lead to the initiator misinterpreting an RNEG and giving the user false information, so reducing the chance of the user being able to find an acceptable set of parameters.

It is therefore recommended that Q should not act on monitor requests when sending an RNEG and that parameters on an RNEG should only include state of transfer, messages and those parameters that contributed to the failure. Each parameter that contributed to the failure should specify a suggested value or range that would have been acceptable or, failing this, the value suggested or implied by P with the operator NE or simply a null value. Furthermore, it is recommended that parameters on a STOP command should obey rules corresponding to those for RNEG.

There are some features of the protocol that it is recommended should not be used, either because they lead to difficulties or because they are erroneous. These are the "parameters on GO" facility, the use of the NE operator on SFT and RPOS commands, the packed form of binary mapping, and data compression on binary files whose word length is not 8 bits.

It is also recommended that the default file type be text. This may seem obvious, but difficulties have arisen because it was not stated in the specification. The problem arises with systems, such as that at CADC, that do not distinguish text files from binary files in their filestore. Such systems need a default when they are faced with a request for a file and the option of choosing either text or binary codes.

Reminders

The reminders are necessary where there is no ambiguity in the specification, but where either some implementors have imagined an ambiguity or where some point is not made strongly enough and some implementors have ignored it. The reminder topics are:

the default for maximum transfer size;
the difference between the Codes attribute and the CS command;

the lack of significance of subrecord boundaries;
the combination of format effectors CR => LF with LF => CR;
and the fact that late errors can result in P and Q having different opinions of the final state of transfer.

Suggested Extensions

There are suggested extensions for:

1. more precise format effectors, including virtual page size;
2. an attribute for code(s) used for storage rather than transmission;
3. a facility for mixed code files, with a change to the CS command;
4. naming of private codes;
5. specifying the ordering of bits within bytes;
6. specifying validity checks on stored data;
7. a STOPREPLY command and the reuse of the old connection;
8. specifying the meaning of FTP records and whether their boundaries should be preserved.

Other Problems

In addition to the problems addressed by the recommendations and suggested extensions, there are some other more fundamental problems with FTP as an interim standard.

Firstly, the protocol is too limited. In particular, it deals adequately with the transfer of complete sequential files but not with file management operations, such as renaming and deletion, nor with the transfer of part files or indexed sequential or random access files. Secondly, the negotiation mechanism is inadequate for achieving near optimal parameters in all but the simplest cases. The solution to these problems will, however, have to wait for a long-term standard File Handling Protocol, but an indication of what this might be like can be gained by looking at the work of BSI DPS 20 WG6. This, however, poses another problem, that of the potential incompatibility between the interim standard and any long-term standard.

Another problem is that there is no interim standard for either an application interface or a user interface, and without these some of the value of having an interim standard will be lost. The application interface will have to include the provision of parameters for the SFT and the receipt of progress and error messages and the final termination state. It would therefore not be difficult to define an abstract application interface, but it will take time to produce a corresponding interface in each of the major programming languages, which is really what is needed. The user interface, designed for humans to drive the initiator, needs to be friendly. Some characteristics that should be included are good help facilities, meaningful progress and error messages, and a way of predefining patterns of parameters so that very few need to be given each time a transfer is requested.

The final problem is that there is currently no centre that can test the conformity to the standard, acceptability or performance of an implementation. Such a centre cannot hope to validate an implementation completely, but it could run standard tests on one and issue an appropriate certificate of acceptance or failure. This would give much greater support to the interim standard than any lists produced by the FTP Implementors' Group and would help assure stability of the standard and compatibility between implementations. It is therefore hoped that the government will establish such a Certification Centre as soon as possible.

File Handling Protocol

So much for the interim standard, but as has been already stated what is really needed is a long-term File Handling Protocol (FHP). Unfortunately, at the current rate of progress in ISO, it will be another 5 years before such a protocol is agreed. Nevertheless, BSI DPS 20 WG6 has already done enough work to give us a feeling of what it should be like.

The FHP should be based on a Filestore Image, which is a system-independent way of viewing files. The Filestore Image [4] should define attributes that files can possess and the range of values for each. Examples of such attributes, other than those in FTP, include the list of users or categories of users allowed to access the file, its date of creation and its organisation (sequential, indexed sequential or random access).

The Filestore Image should also define the operations that can be performed upon the filestore. For some of these there may be additional attributes defined, relating to the specific activity. For instance, file transfer operations should have associated transfer attributes, as in FTP. Other possible operations include selecting or deleting a file and reading or changing file attributes (e.g. to rename a file).

In addition to operations on single files, there should also be operations on sets of files. These should enable files to be grouped for ease of access. This grouping might correspond to the directory structure found in many real filing systems.

There should also be operations on part-files. These could be defined around a single subunit of a file, called a record. This would be sufficient to allow keyed access to records in indexed sequential or random access files.

Conclusions

There is a need for a single internationally accepted standard File Handling Protocol to enable any filing system to be accessed remotely in a system-independent manner. Unfortunately, such a standard is probably 5 years away.

There is, however, an interim standard File Transfer Protocol that can be used on PSS. It is more limited in concept, but has already proved on EPSS to be very good for the transfer of complete sequential text files. This interim standard is currently maintained by the DCPU FTP Implementors' Group, which has issued lists of recommendations and reminders to implementors to clarify the specification. It has also produced a list of suggested extensions, which can be used as input to the long-term work of BSI DPS 20 WG6.

Nevertheless, there are still some fundamental problems requiring urgent solution. Firstly, there are no corresponding interim standards for the interfaces to be presented to application programs or human users. Secondly, there is currently no centre that can test an implementation of the interim standard for conformity, acceptability or performance. This is not only a political problem, but also a technical one because, as yet, we do not know how to test implementations of such a complex protocol.

References

1. A Network Independent File Transfer Protocol, prepared by the UK High Level Protocols Group, HLP/CP(78)1, December 1977, available from Dr P.F. Linington, Data Communication Protocols Unit, Computer Laboratory, Corn Exchange Street, Cambridge CB2 3QG.

2. Recommendations and Reminders to Implementors, produced by the Data Communication Protocols Unit FTP Implementors' Group, edited by D. Rayner, July 1979, available from Dr P.F. Linington, address as in [1].

3. Suggested Extensions of FTP, produced by the Data Communication Protocols Unit FTP Implementors' Group, edited by D. Rayner, July 1979, available from Dr P.F. Linington, address as in [1].

EXPERIENCE IN THE DEVELOPMENT AND USE OF NETWORK CODE ON SOME GEC 4000 COMPUTERS

PAUL BRYANT

Rutherford Laboratory

SCIENCE RESEARCH COUNCIL

RUTHERFORD LABORATORY
COMPUTING DIVISION

EXPERIENCE IN THE DEVELOPMENT AND USE OF NETWORK CODE ON SOME GEC 4000 COMPUTERS

P Bryant

24 September 1979

During the last year the Interactive Computing Facility of the Science Research Council based at the Rutherford Laboratory has placed GEC 4000 series Multi User Minis at Bristol, Cambridge, Glasgow and Newcastle universities as well as Cranfield Institute of Technology. Further installations will be at Bradford, Cardiff and Birmingham universities. The computers are used principally for the support of SRC-sponsored research in a variety of disciplines.

The computers are supported from the Rutherford Laboratory and all operating systems and basic software is supplied from Rutherford and delivered to the sites via the SRC network. This is to avoid the need for systems programmers at each site with the attendant costs in terms of manpower and system testing time. User support is provided from Rutherford. Much of this support is provided via the network by 'conversations' over the network and by the computer 'POST' system.

The details of the SRC network were reported to Networkshop by John Burren. The network is based on X25 and connects together many of the SRC computers. The GEC computers support 3 high level protocols. The ex-EPSS ITP protocol is used for interactive traffic. FTP-B is offered and HASP protocol is used for job transfer to the IBM computer.

The block diagram shows the structure of the GEC network code developed at the Rutherford Laboratory. The boxes represent operating system processes and the sizes relate to the size of the processes including data spaces.

Interactive facilities are provided by two processes - one for user ITP and one for server ITP. In the case of user ITP the user does not need to log in to the local GEC COMPUTER. This makes user ITP very cheap to run (less than 1% of the computer) whereas if a user had to be logged in, the overheads would be far higher. The local terminal control sequences are disabled, such as ones which affect character conversion or break in, this prevents any ambiguity as to which computer is absorbing the user utterances. In fact the only commands actioned locally are:

@Q this cancels the call
@Y this causes a break in and sinks any data in transit

Note that backspace is actioned by the host. This was found essential due to the GEC facility which allows a line to be 'cleaned up' by the host and output on the terminal for reentry after possible reediting. In this case the user could backspace a non-existent record. Clearly better ITP facilities are required in this area.

Inactivity of ITP calls for 11 minutes causes User ITP to break the call. This time out is slightly longer than the automatic log out facility so that the automatic logout occurs before call disconnection. Logging out does not break a call. However call disconnection for any reason does cause an automatic log out. Ideally it should not do this but allow a user to reconnect to a session. This is difficult but is a possible enhancement.

The user accesses the ITP system by typing !!ITP followed by the host 'title' or alternatively the full address of the host. In the case of a title, this is translated to the full network address by means of a file of titles. SRC have developed a naming scheme for hosts. The first two characters represent the site (based on the ex ATLAS ex 1906A and current ICF user names) the next letter is the machine type, the next is a letter to remove ambiguity and the last the protocol. For example GWGAI represents Glasgow GEC, machine A and ITP. Note that the user can always force any address sequence through by explicitly giving it. Thus the user can access all possible hosts with any protocol (although this may not always be meaningful).

Server ITP interfaces to OS4000 via the 'Line Terminator' processes which normally interface to the multiplexer code. A group of LTs (typtically 3) are dedicated to ITP although a desirable modification is to pool the LTs used for local connections and ITP connections.

A serious shortcoming is that the operator cannot control the maximum number of clients on user ITP connections and is unable to find out how many user ITP calls exist or communicate with those users. This will be implemented as soon as possible.

Interestingly the implementation of ITP (both user and server) was found to be very simple and only absorbed a couple of weeks of effort.

FTP is an order of magnitude more complicated. It is provided by two very large processes - FTP and MAFS. FTP is effectively the analysis of data according to the 'blue book'. MAFS is FTP-B independent and deals with access to the filestore, spooling of transfers, queueing of tranfers and user control of file transfers.

The FTP process deals with a fairly option-rich version of FTP-B and is strictly FTP-B. For Q process, FTP analysis the FTP-B parameters and then requests MAFS to verify the parameters. If the parameters are acceptable, MAFS connects FTP to the filestore (or device spooler) whence the transfer proceeds independently of MAFS. At the end of the transfer

MAFS again is informed. MAFS then informs the filestore owner of the success or otherwise of the transfer. If logged in, a message is sent to his terminal; otherwise a message is 'posted' to him.

FTP using P process is rather more complicated. The user first invokes the TRANSFER command which asks for the file transfer parameters. They are syntactically checked. The host 'title' is translated from the same file as used for ITP. Again the user can force through an explicit address. MAFS is then sent the parameters and the parameters are checked semantically and if correct they are placed in a non-sequential queue. Non sequential, as it is important that a blocked file (ie, the host is not available) must not block files that can be transferred. For outward transfer the file is copied to a spool After transfer, successful or otherwise, the user is informed. Note that unsuccessful transfers are reported (eg, password error) and thus a user knows of unauthorised attempts to access his filestore. It is important to note that the file transfer scheme maintains the integrity of the GEC filestore.

The TRANSFER command also allows the user to examine the FTP queues. He can also delete or abort transfers he has initiated.

There is currently no mechanism for automatically purging the queue of files that cannot ever be transferred, ie, the host does not exist. It is proposed that the user can demand that a file must be transferred by a given time date - if not the transfer is deleted. A default of a day or so would be reasonable.

Originally MAFS started as many transfers as it could - maybe all to the same host. It was discovered (painfully) that it is better to start only one transfer per host as some files then become available sooner. If only one transfer is allowed per host in general, then some priority mechanism is required. This is under consideration.

As well as TRANSFER there is a simpler version which allows network POST. This allows amazingly easy use of the network for 'electronic mail' if one can use such an overworked term.

The HASP system is provided by two processes. The MAFS system is used for returned data. This allows data to be returned to the filestore or lineprinter. Data to the IBM machine is via a standard GEC spooler.

MAFS is not only used for the X25 network. MAFS can 'loop back'. This allows files to be moved between users file spaces. Otherwise this is difficult due to the protection of the filestore system. In principle MAFS should not loop but this should drop down to transport level and back again. This would be less efficient with absolutely no benefits.

A simple asynchronous single call-call and link level is available. This is used for file transfer to the PRIME. This protocol is also available on PDP11 under RT11 and RSX11 (including FTP-B). Versions on M6800 and 8080 are under development. These will be used for data collection and word processors. The protocol complete with FTP-B (a very option poor version) takes 2 Kbytes on a micro. The system has proved very reliable and useful. ITP has been available since Janury 1979 and FTP since July 1979. Recently a complete operating system and system software

was transferred from Rutherford to Bristol over a 2.4 K line. It was done over a weekend. It involved 300 files and took 12 hours. The machines were unattended and were providing a service during the process. The operation was a success.

The complete network code takes about 81 Kbytes of code, data and buffers. It took some 2 man years to write over a period of 9 months. It turned out to be easier than expected and more bug free ~~than~~ expected at this early stage. The overheads of use have not been measured but experience suggests that the system has little or no effect on interactive performance.

From the implementation point of view FTP and transport level have some problems. FTP is option rich - too rich. Much of the complexity of the code, and hence its size, is due to these options, many of which will probably never be used. For example, it would have been better to restrict the character code to ISO. A single carriage control convention. Only 8 bit binary data. In the case of machines needing other standards, my view is that they should provide a mapping which would not seem difficult.

Transport level is implementation-wise curious. It is neither a separate identifiable level nor is it purely an interface definition. In the GEC implementation it turns out to be mainly an interface definition. This seems to meet the philosophical aims of transport since other call levels can be supported without altering high level protocols and vice-versa. The address features of Transport which can be regarded as an interface definition splurges through the high levels. For example, some of it is embedded in the file of 'titles'. Blocking of FTP data into the 63 bytes blocklets is defined in FTP-B. This is unfortunate. The implementor seeing he has to do blocking and deblocking realises he has to do some buffering. He thus asks what the optimum buffer size is. miraculously he learns that the best buffer size is the X25 packet size and he uses it. This goes away from the spirit of Transport. My view is that the blocking deblocking should be a feature of Transport from the implementation point of view.

Nontheless, these points are not too serious and the whole system is aesthetically as well as practically successful.

A last point. It would be a disaster if JTP were not based on FTP. Another 21Kbytes of JTP corresponding to FTP would be a considerable burden to write and maintain.

I would like to pay tribute to the implementors : Dave Toll, Andrew Dunn, Tony Salter and Graham Robinson who have produced a system which will be hard to improve upon.

P E Bryant
24 September 1979.

GEC 4000 SERIES NETWORK CODE

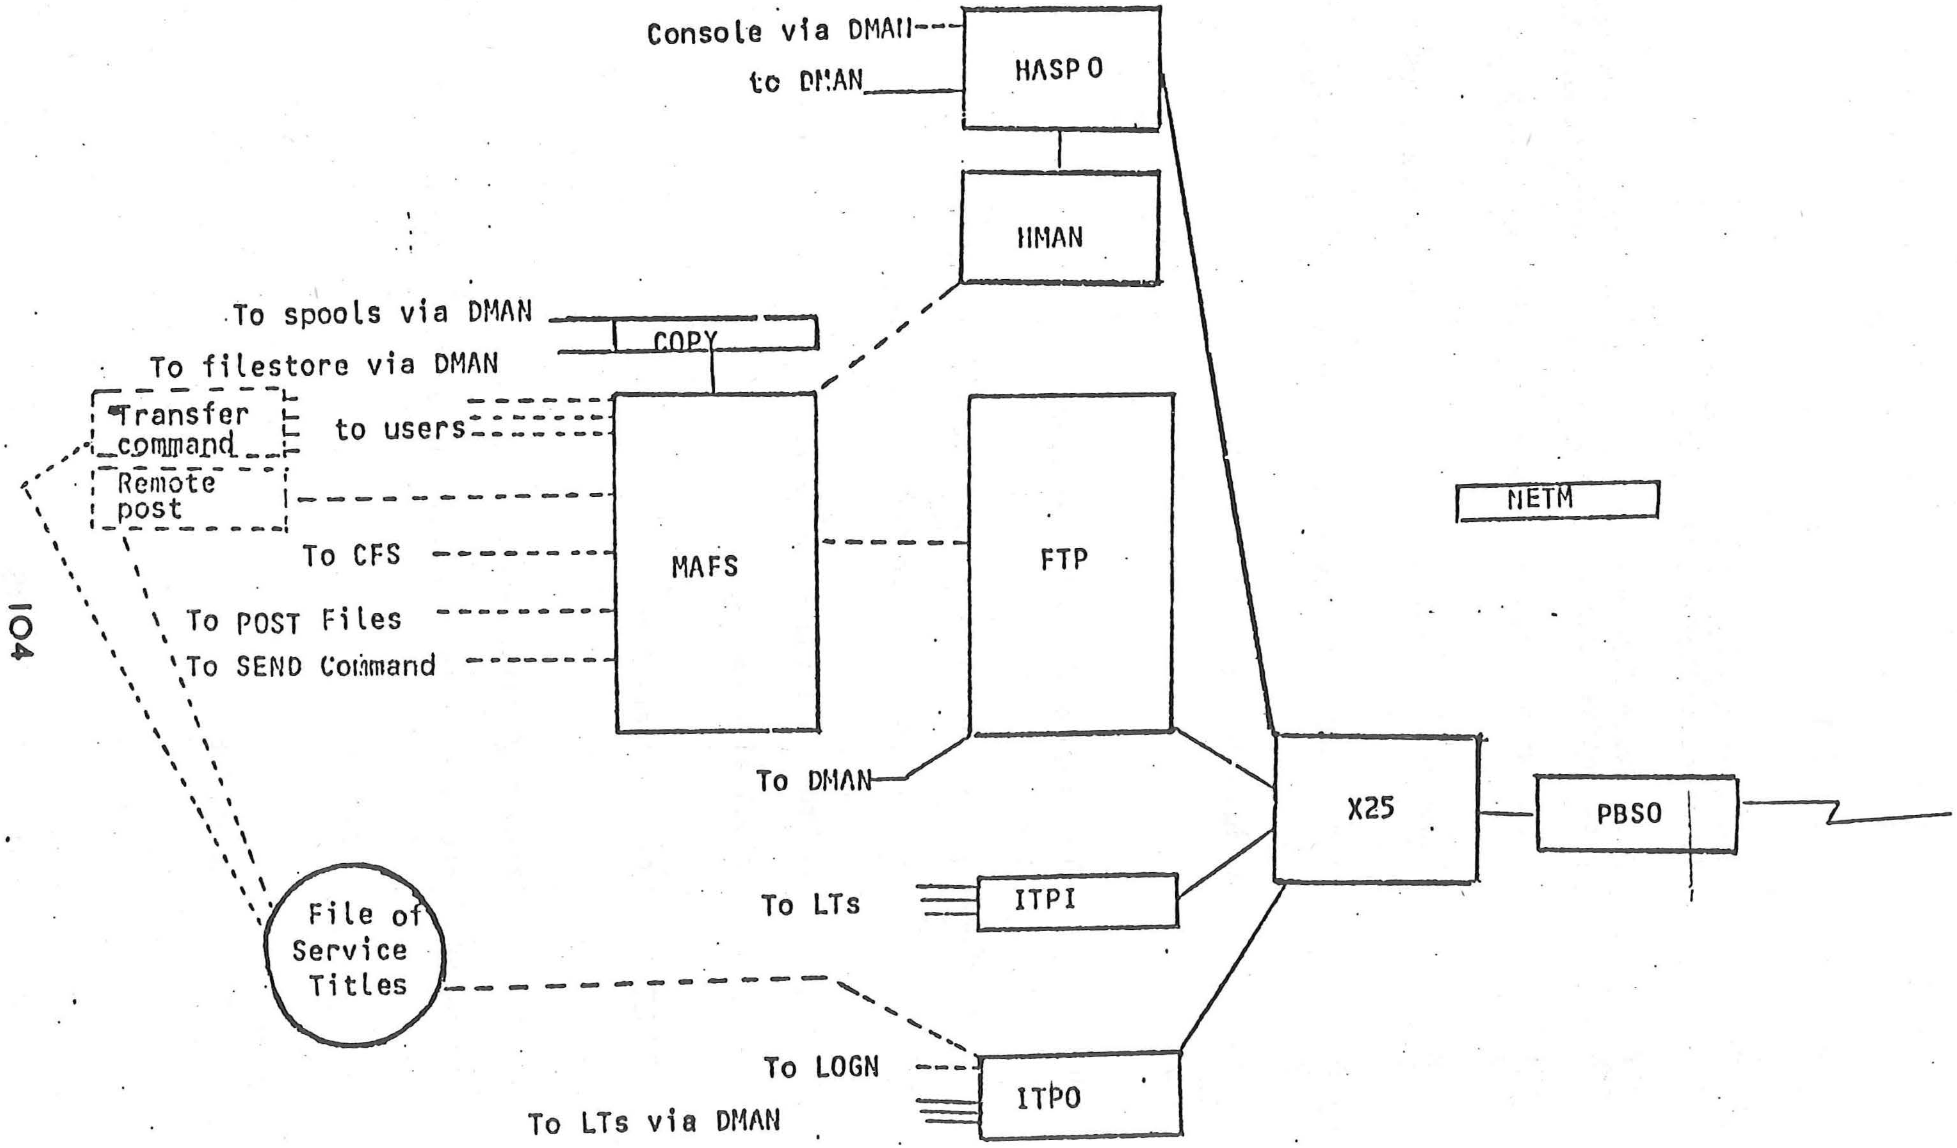

NOTES

PBSO 4.5K fixed	Contains link level code Can have further processes PBS1 etc.
X25 10K	X25 level 3
HASPO 7K	For each HASP connection a HASP process is needed, HASPO, HASP1, etc
FTP 21K	Deals with file transfer protocol-B. It can connect to alternative call level protocols, ie, asynchronous link to PRIME
ITP1 3.5K	Deals with interactive terminal protocol into OS4000
ITPO 3.75K	Deals with interactive terminal protocol out of OS4000
MAFS 21K	Deals with file transfer control to and from OS4000
NETM) 1K) HMAN) 7K)	Module Managers
COPY 2K	For copying files to the spool
Transfer command	To enable users to initiate transfers
Remote Post	To post messages to remote users

Key

---- Control & Information Routes

Principal Data Routes

Group of Data Routes

Group of Control Routes

The size of the boxes relate to the data and code sizes

A Process

A user job

A file

REPLACEMENT OF THE NPL DATA COMMUNICATION NETWORK

KEITH BARTLETT

Data Communication Protocols Unit

Replacement of the NPL Data Communication Network

by K.A.Bartlett

The Present Network

The NPL Data Communications Network (DCN) has been operational in its present form since 1973 during which time its has become a site service alongside other essentials such as water and electricity. It is a common carrier system and has allowed the independent development of a wide range of remote access facilities serving both the scientific and administrative populations.

The network currently supports computational services, databases, word processing, message services and special-purpose scientific experiments. There are approximately 200 subscribers including 25 computers and about 120 visual display units. It is available day and night carrying between 500,000 and 1,000,000 packets every 24 hours

The access links connecting subscribers to the single central switch vary in capacity but they can work up to 10k characters per second. It is actually a pure datagram network but a simple virtual call protocol was necessary for terminal handling and this is used by all subscribers. The single packet-switch can handle a maximum of 250 to 500 packets per second depending upon the packet size. A 'Terminal Processor' (or PAD, in current terminology) is incorporated in the same computer as the packet switch and can handle up to 250 data blocks per second - again depending upon size.

The design of the network commenced in 1968 as part of NPL's original work on packet-switching as a technique for data communications. The central switching computer is a 32k Honeywell H516 purchased in that year and special purpose hardware was built to connect subscribers to this single switch. The experimental purpose of the network was rapidly swamped by its use as a service and the present bottleneck of connections was reached in 1976.

The hardware is still functioning well but if demands for further connection to the network are to be met, more will have to be built and it does not seem sensible to propagate this early and rather expensive design. In 1978, a decision was made that a replacement network based on X25 should be obtained from commercial sources and the resultant specification against which the new network will be purchased was finalised early this year.

The present network has a gateway to EPSS and this is a complex function. The only connection to PSS will be made through the new network which, as an X25 system, should present few problems in connecting to the Post Office service.

The New Network

The replacement network will be X25 compatible. Initially, it will co-exist with the current network with a gateway connection between the two. The decision that the new network would offer only X25 connection (plus a PAD) was taken entirely on the basis that computers and terminals for connection to such a network would be readily obtainable as standard manufacturer supported products. All connections to the current network have expensive custom-built software and hardware and the same problem will occur in the next few years with any form of ring local network. The new NPL network will be a small-scale version of PTT networks and the wide range of products designed to work with these will be usable in the Laboratory without further interfacing.

The final form and capacity of the network is difficult to determine at this stage but will probably be about 100 computers, 500 interactive terminals and 100 other devices. These will be from a variety of manufacturers as the requirements of the users will differ. It will save a large of amount time and money both in development and procurement procedures if the specification to which these products are supplied is one which is widely used and manufacturer supported.

The procurement procedure is under way and it was originally hoped to have the new network in service in time for connection to PSS on day 1 of that service. However, there have been a number of delays and it is now unlikely that this will happen.

Specification

Good data communication systems are addictive and because it not possible to predict the final form and layout of the network, it is specified in terms of data exchanges. The initial purchase will probably be for two or three of these. To allow the maximum flexibility in configuring the data communication requirements of the site, each exchange represents a complete network - for which it is the sole switching point. The exchanges must therefore be interconnected by gateways but since each network is identical, the actual gateway problem is very small. There are several advantages to this scheme including flexibility and security and this applies particularly to the Laboratory connection to PSS or Euronet. These services can be connected to any point on the Laboratory network without requiring a unique external service gateway.

One disadvantage of the scheme also arises from considering how to connect to a PTT service. The PTT will not accept the NPL (or any other non-PTT system) as a network i.e it will not interconnect using X75. The private network must appear to the PTT as a multiple call DTE via an X25 connection. This means that the NPL connection to PSS will be X25 and for flexibility and compatibility therefore, all the internal internetwork connections (i.e. between the data exchanges) must also be X25 (not X75). This is not a serious disadvantage and certainly saves any private network agency worrying about X75.

The Data Exchanges are specified in a modular manner. Major components are an X25 packet switch and a PAD - which appears to the packet switch as a multiple call DTE thus allowing the simple connection of further standalone PADs if required. Other major functional items are the gateways (major in concept but trivial in practice) and a Network Management Centre (NMC).

'Access Units' are also specified. These are the means by which the subscribers connect to the exchanges and are sub-divided into Packet Mode Access Units and Character Mode Access Units. An obvious example of a Packet Mode Access Unit is one which provides LAPB as level 2 of X25 although other implementations of level 2 will also be required including one link which must work at a clock rate of 1 megabit/second. The most common access unit for character mode devices will be that which gives 20ma current loop signalling.

An exchange is specified by calling for a number of modules which would typically be a packet switch, a PAD, a Network Management centre and a number of access units of different types.

The packet switch is required to complete a call attempt in 50 milliseconds and to forward subsequent packets in an existing call within 10 milliseconds. A call must be cleared within 50 milliseconds and this time includes updating charging and statistical records - as does the call attempt time.

These performance figures apply to all exchanges but the actual throughput and call capacity of an exchange will be specified at time of ordering. The first procurement will be for two slightly different exchanges. The larger of the two will handle an average of 150 packets per second with a peak of 400 packets per second. It will support 150 simultaneous calls with a call creation (or breakdown) rate of 10 calls per second. This exchange must support ten X25 access links, one at 1 Mb/s, six at 48kb/s and three at 19.6kb/s.

The PAD will respond to a control character within 50 milliseconds and echo characters (when required) within 10 milliseconds. Again, the PAD configuration can be specified as required but the first one will be for 30 terminals, 20 of which will use current loop signalling.
The PAD will work according to X3,X28,X29 as recommended by the SG3 working party.

The reliability figures are not particularly stringent as we believe that the end-to-end procedures must be prepared to bridge short term network failures - especially where human users and long set-up or logon sequences are involved. A packet switch must not suffer a complete failure more than once in every 500 hours. The total out-of-service time for any subscriber due to exchange faults must be less than 1 hour in every 2000. Packet loss by the switch is specified as better than 1 in 1,000,000,000.

The Network Management Centre may not be within the exchange but can exist within any nominated computer on the network. There must, therefore, be well-defined interfaces to a number of modules within the exchange which can offer logging, statistical control and test facilities to the NMC which will be able to perform a number of functions including:-

Control and configure the network
Allocate mnemonic addresse to subscribers
Enable or disable a subscriber
Obtain the service and call status of any subscriber
Obtain network and subscriber statistics (24 hour period)

It will also allow a number of tests to be performed by subscribers or a network operator.

Conclusion

The facilities in the exchange and the performance levels specified are such that it should be easily obtainable from a number of suppliers and be a suitable component for a wide range of networks. The restriction to PTT recommended modes of working - X25, X3 etc. will make it compatible with a large number of commercially supplied products without the need for special engineering. All these factors will keep the real costs of installing, running and using the network to a minimum.

A MAINFRAME-TO-ETHERNET CONNECTION

PETER GIRARD

Rutherford Laboratory

A MAINFRAME TO ETHERNET CONNECTION

1. Introduction

The Rutherford Laboratory has been involved in the design and construction of Ethernet hardware for several years, and this has now reached the stage where it can be given a real function to perform in a production environment.

The Laboratory does not have a "Local Area Network" in the currently understood sense of that term. It has had for many years a continually expanding star network of HASP workstations connected into the 360/195 mainframe, and a large number of interactive terminals also connected to the mainframe directly or through the workstations. More recently, in collaboration with Daresbury Laboratory, an X25 network has been set up. This too has a port into the 360, and an increasing number of workstations and terminals now use the 360 by this route.

The Ethernet work has been going on separately and in parallel with the above developments to the production system. It has given some people at the Laboratory the opportunity to acquire knowledge and expertise in this field, which will undoubtedly be valuable, if at some stage a Local Area Network is set up. At present, however, there are no definite plans to do this.

The method chosen for giving the existing Ethernet hardware a useful function in the production system is to incorporate it into the SRC Network, by connecting it between some of the latter's DTEs and DCEs. Figures 1 and 2 illustrate this in two different ways, and from fig.2 it can be seen that the function of Ethernet will be very similar to that of an SRC Network DCE. As such, it is expected to relieve the main network, and the 4080 packet switch in particular, of a substantial amount of traffic.

2. The Ethernet Hardware

The central feature is a single coaxial cable, functioning as a passive broadcast medium. This is called the Ether. Equipment in the form of Ethernet "nodes" is connected to it simply by tapping into the cable. Four such nodes have actually been built, as shown in fig. 3. The nodes, in turn, are each connected

to a "host", and in practice these will be the SRC Network machines shown in fig. 1.

When a node has data to transmit, it waits until the Ether is quiet, then broadcasts the data packet. This has an Ethernet header on the front (see fig.4) starting with the address of the destination node; all other nodes will detect but ignore the packet.

If two or more nodes begin to transmit at roughly the same time, all the packets involved will be corrupted and must be transmitted again. A sender can detect interference by listening to its own transmissions; if a collision has occurred, it jams the Ether deliberately for long enough to ensure that all the nodes are aware of it, then waits for a short 'random' interval and tries again.

Each node contains a Motorola M6800 microprocessor and 16K of memory.

3. Ethernet Software

If a packet reaches the receiving node and its CRC is correct, the receiver will send back an acknowledgement packet. If the sender fails to receive a correct acknowledgement, it will timeout and re-transmit the packet. The sequence number in the Ethernet packet header ensures that the receiver can detect and discard any duplicate packets.

Only one acknowledgement can be outstanding from a particular node, but a node can have acknowledgements pending from several different nodes at the same time. Thus a delay or a fault in one node will not interfere with the flow to other nodes.

4. Incorporation into SRC Network

To function as a normal SRC Network packet switch (DCE), Ethernet would require X25 Level 3 software in its nodes. To avoid implementing this, it was decided to treat Ethernet strictly as a Level 2 entity: It should look like a Level 2 'link' between any two SRC Network machines to which it may be connected.

It is of course unlike a conventional link, to the extent that it is connected to multiple hosts rather than

just two. However, the concept can be maintained by considering each pair of hosts (when actively supporting calls) to be connected by a separate logical link, even though all the logical links to a given host actually share the same physical medium.

The Ethernet line drivers in the hosts must put their X25 packets into the data field of Ethernet packets before transmitting them. Used in this way, the Ethernet is of course completely transparent to X25 Level 3 and all higher levels of protocol.

5. Impact on X25 Level 3

Treatment of Ethernet as a set of Level 2 logical links has several minor repercussions on the X25 Level 3 implementation in the host machines:-

1) Two DTEs can face one another across Ethernet. This is not a problem except that "call collision" is more likely to happen than with a normal DTE/DCE interface.

2) Each logical link through Ethernet must be considered independent for LCN-assignment purposes. Thus the Ethernet driver in a host must be able to tolerate the same LCN on different logical links through Ethernet.

3) Each host must know the Ethernet address of all nodes with which it may communicate. This address also serves as the most convenient way of identifying a logical link.

4) Whereas a link normally consists of one hop, crossing Ethernet involves 3 hops; hence a larger X25 window is likely to be needed than on a conventional link.

6. Physical Connections to Hosts

Only the connections to the 4080 packet switch and the IBM mainframe have been worked out in detail.

The connection to the 4080 is almost complete. It uses a store port on a 4080 storage module, enabling the node to write directly to and read directly from the 4080's memory. This is supplemented by a connection to an asynchronous board on the 4080, through which control

signals that will cause interrupts can be passed in either direction.

The connection to the IBM mainframe will be through an existing Parallel Data Adapter on an IBM 2701 Data Adapter Unit. The 2701 connects directly to an IBM channel. This mode of connection should allow much higher throughput than serial communication using Binary Synchronous, and it was hoped that a particularly simple protocol for reading and writing could be used, avoiding a regular handshake. Unfortunately, this does not look like being feasible without a lot of work, and it is likely that a simple handshaking technique will be used after all, at least in the first instance.

7. Summary

Our motive for connecting the Ethernet hardware to the SRC Network is to demonstrate and use it in a production environment, where it can perform a useful function. To do this, it is being connected between several of the local SRC Network DCEs and DTEs. The simplest possible form of connection has been chosen, in which Ethernet is transparent to X25 and all levels above it.

If Ethernet were being used to support a Local Area Network, it is unlikely that it would be connected in a similar way. To achieve sufficient flexibility in the design and use of such a network, it would almost certainly be necessary to treat it as a network in its own right, and connect it to other networks through Transport level gateways.

Fig 1

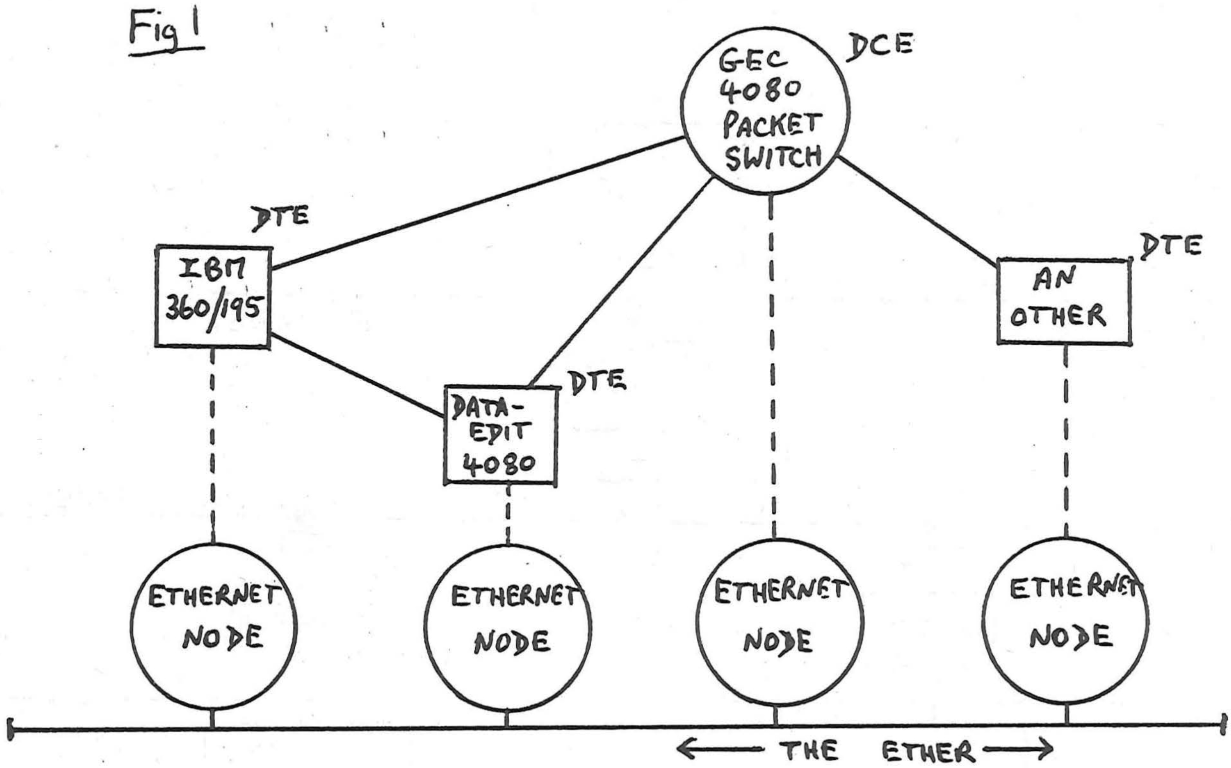
GEC
4080
PACKET
SWITCH
DCE
DTE
IBM
360/195
DTE
AN
OTHER
DTE
DATA-
EDIT
4080
ETHERNET
NODE
ETHERNET
NODE
ETHERNET
NODE
ETHERNET
NODE
THE ETHER

Fig 2

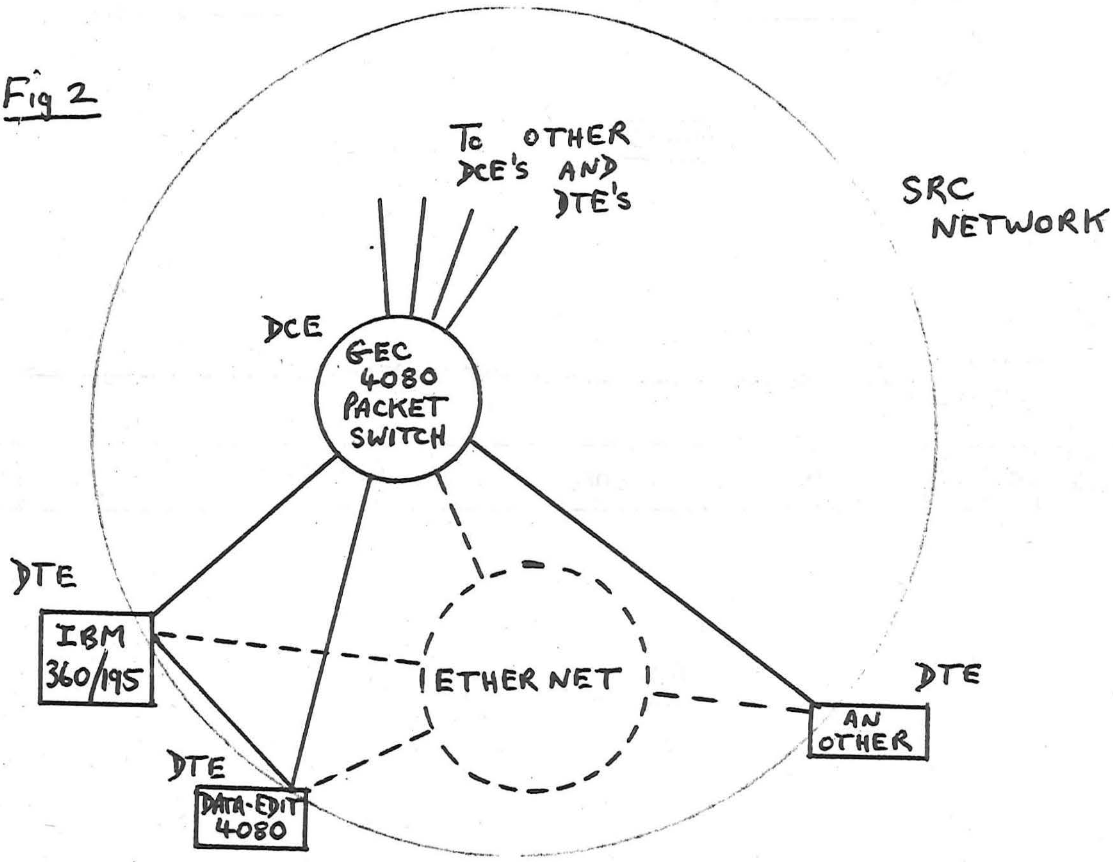
To OTHER
DCE'S AND
DTE'S
SRC
NETWORK
DCE
GEC
4080
PACKET
SWITCH
DTE
IBM
360/195
ETHERNET
DTE
AN
OTHER
DTE
DATA-EDIT
4080

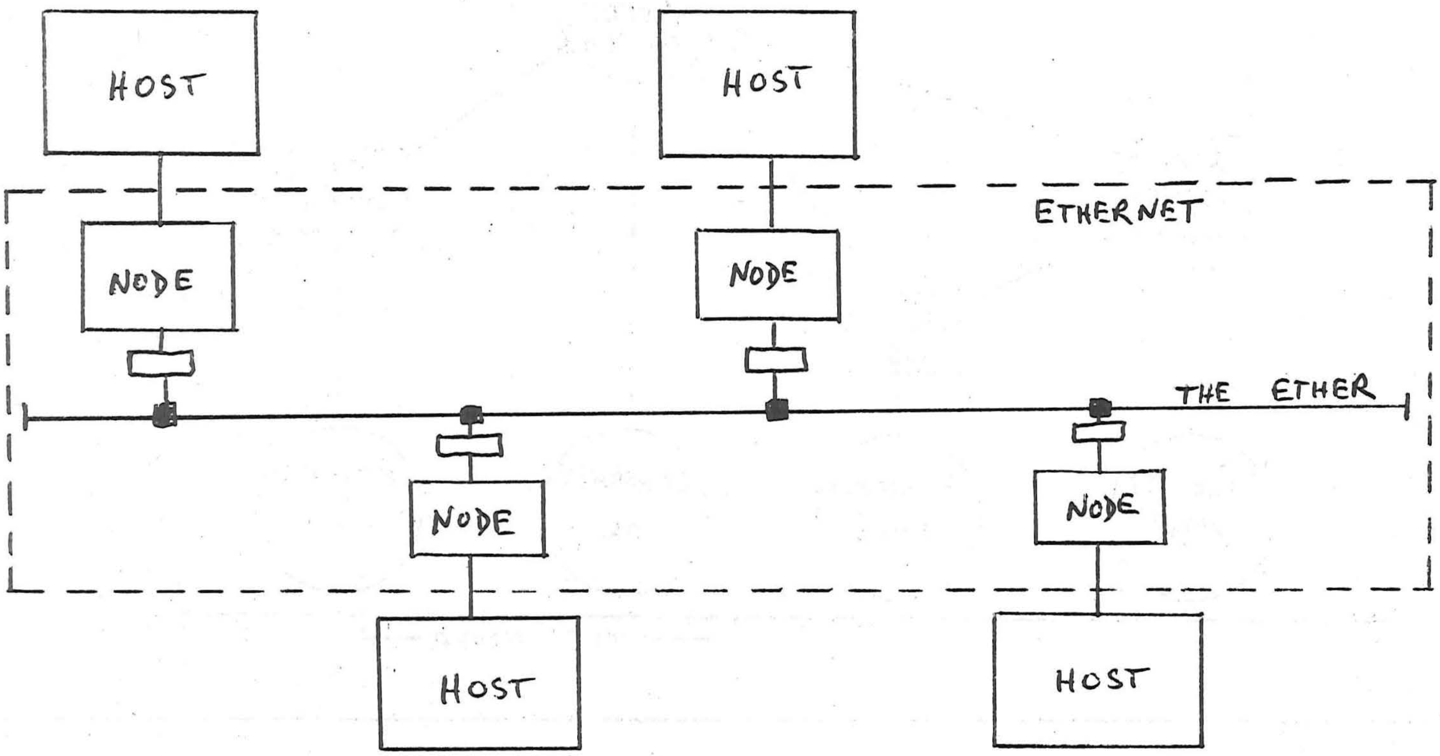
HOST
HOST
ETHERNET
NODE
NODE
THE ETHER
NODE
NODE
HOST
HOST

Fig 3

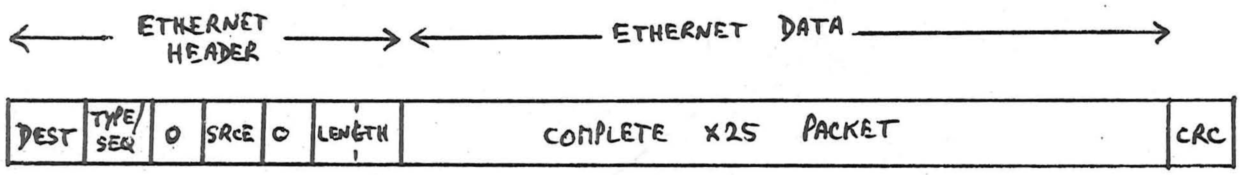
ETHERNET HEADER
ETHERNET DATA
DEST
TYPE/SEQ
0
SRCE
0
LENGTH
COMPLETE X25 PACKET
CRC

Fig 4

LOW LEVEL PROTOCOLS FOR THE CAMBRIDGE RING

STEVE WILBUR

University College, London

Low Level Protocols for the Cambridge Ring

Steve Wilbur

University College London

Abstract

This paper discussed some pitfalls encountered in designing low level protocols for the Cambridge Ring. Some possible solutions are suggested, bearing in mind the problems posed by internetwork bridges.

Presented at "Networkshop 5",
University of Kent at Canterbury.
19 - 21 September 1979

1. Introduction

At University College London in the Department of Statistics and Computer Science, we have built a Cambridge Ring. It is identical to the original at Cambridge University except that printed circuit boards have been designed in order to speed production. The ring is at present connected to two PDP-11/10s and a PDP-11/34 in the Teaching Laboratory and to one or more LSI-11s in the Communications Reseach Laboratory. Shortly we hope to have terminal multiplexors also attached to the ring, and to change the existing programmed transfer interfaces for intelligent DMA interfaces.

In the Teaching Laboratory we expect to be using the ring for:

- sharing peripherals
- terminal switching
- distributed operating system

In the Research Laboratory possible applications are:

- interconnection of protocol convertors
- internetwork experiments
- local area network protocol requirements

2. Cambridge Ring

The Cambridge Ring consists of a set of repeaters connected one to the next in a circular fashion by two twisted pairs (four wires). The propagation delay in the wire and repeaters is used to store data on the ring, and at 10Mhz clock rate about 4½ bits are stored/100m of cable; 1½-2 bits are stored in each repeater. A mini-packet (m-p) (described below) is 38 bits long, and if the delays in the natural ring are insufficient a shift register is inserted at the so-called "monitor station" to pad the length to 38m + n bits, where m is the number of m-ps circulating and n is the number of "gap" bits (>1).

Thus, with the above arrangements data circulates round the ring at 10M bps serially. Attached to the repeaters is a station unit which converts parallel data from/to a device or machine, to/from the serial data on the ring. The station knows about mini-packet structure and addressing. The m-p format is

shown in figure 2. Each m-p holds 6 status or maintenance bits, 8 destination address bits, 8 source station address bits and 16 data bits. Of the status/maintenance bits, the full/empty bit at the head of the m-p is used to indicate whether the m-p is free. When data is inserted into a m-p this bit is set full and the m-p moves to its destination. On arrival a copy is made of the data and origin and the trailing status bits are modified in the ongoing m-p. When the m-p arrives back at the source station the full/empty bit is set empty and data is checked on a bit-by-bit basis; any error is reported to the machine/device interface. If no errors occurred, the status bits inserted by the receiver are presented to the machine/device interface. Because the originator cannot immediately re-use the m-p and has to pass the m-p downstream a round-robin scheme exists for access to the ring.

The receiving logic of the station has an 8-bit register, called a "source select register" (SSR). When the value 0 is set in it all m-ps addressed to the station will be refused, when set to 377_8 all m-ps sent to it will be accepted (provided previous m-p has been processed), and when set to any other value ($1\text{-}376_8$) only m-ps from that station will be accepted. The SSR may be changed from the machine/device interface.

Four possible status values may be returned to the originator of an m-p. They are:

- accepted
- busy - last m-p not yet processed
- not known - no such station
- unselected - station exists but SSR blocking

There is a parity bit in each m-p. This is checked and set by each station in turn around the ring. If an error is detected a flag is set and a fault packet is sent to address 0 in the next empty packet. Indications of parity error are never passed to the machine/device interface, thus the parity bit is purely for maintenance purposes.

It is also worth noting that

- the receiver has no check on data validity at the hardware level
- only the transmitter knows of possible data errors at this level, because of the bit-by-bit comparison.

3. Low Level Protocol

A study of our proposed applications led us to believe that the primitive communications mechanism we wanted was a process-to-process message handler. This we tried to do with the datagram protocol of Figure 3, where Dest-P and Src-P are process identifiers, Ack-H is the header acknowledgement and Ack-D is the data acknowledgement. It is envisaged that it would be used in the environment of Figure 4, where there are a queue of messages awaiting dispatch and a queue of requests for incoming messages at the remote site. If there is not a request queue entry corresponding to an incoming header then Ack-H is negative. It is also negative if the receiving buffer is too small, but this time a parameter indicates the buffer size. Thus this header acknowledgement provides a flow control mechanism. If a negative header acknowledgement is received, the data is not sent.

4. Half Duplex Algorithm

It is possible to write a half duplex algorithm to deal with this protocol.

TRANSMIT REQUEST	RECEIVE REQUEST
. Wait if already receiving	. Wait if already transmitting
. Set SSR=0	. Set SSR=SRC or SSR=377_8
. Send header	. Wait for m-p
. Set SSR=DST	. If SSR=377_8, set SSR=SRC
. Wait for Ack-H	. Wait for rest of header
. Send data if OK	
. Wait for Ack-D	. Acknowledge
. Set SSR=0	. Wait for data
	. Acknowledge
	. Set SSR=0

Initially SSR=0

5. Full Duplex Algorithm

A possible full duplex algorithm could be simpler than the half duplex one, eg.

TRANSMIT COMMAND	RECEIVE COMMAND
(Assume SSR=377_8 or SSR=SRC if receive in progress)	. Set SSR=377_8
	. Wait for m-p
	. Set SSR=SRC
. Send header	. Wait for rest of header
. Wait for Ack-H	. Send Ack-H
. Send data	. Wait for data
. Wait for Ack-D	. Set SSR=377_8
	. Send Ack-D

There are, however, two problems in the full duplex case. The first occurs when a block is being sent from A to B and another one is being sent from B to A simultaneously. It is obviously impossible to distinguish whether m-ps arriving at B are data from A or acknowledgements. There are two solutions to this:

1. Separate data and acknowledgements in time, ie. delay sending ack until all data sent
2. Add extra marker bit in m-p to distinguish data and acks. (In the LSI version of the ring extra bits for this purpose are being included).

The second problem occurs when three or more nodes are engaged in cyclic transfers. If A sends to B, B to C, and C to A simultaneously, then deadlock can occur.

Simultaneity can arise in two ways:

. if A, B, and C all start transfer on same cycle of ring with 3 m-ps in it.

. all nodes are in header state simultaneously.

Deadlock occurs because once the SSR is set to other than 377_8 (or 0) only one channel into a node exists. In this case A can only receive data from C and acknowledgements from B are blocked, and similarly for other nodes, see Figure 4.

The following software solutions exist:

1. Constrain to half duplex in header sequence.

   Possible but tedious. Slight difficulty comes in changing SSR from 377_8 to 0 without missing incoming header

2. Treat data and header as two blocks

   Potentially results in queue of half started transfers

3. Forget about Ack-H, combine it with Ack-D

   Results in waste of:

   Bandwidth - not very important in local nets

   Machine cycles - processing of data by crude interfaces. Intelligent DMA interface would eliminate this.

4. Use "Refused" status bit as NACK.

   By changing the SSR to 0 when receiver wishes to refuse a header, a negative acknowledgement can be passed back but without any parameters (reason). There is a further problem in that the receiver does not know when it has refused the next incoming packet - a timeout is too gross unless your machine has very fast clock (100μS).

Two hardware solutions are suggested:

1. Provide secondary input channel.

   By providing a second SSR which is used when the previously proposed data/control indicator is set to control, one essentially is provided with a secondary input channel and explicit ACKs can be sent without problem.

2. False refusal.

   If it were possible to set a flag in the receiver logic which caused all m-ps addressed to the receiver to be marked refused irrespective of whether they matched the SSR or not, then the originator could treat this as a NACK. If no other change were made to the receive logic the receiver would get an interrupt when the NACKED m-p arrived and so gets positive knowledge of when the NACK was sent.

With the implicit acknowledgement of 2, there are problems when a gateway to another ring is traversed. Unless the NACK is explicit, it will only travel one hop back for every m-p sent. Therefore the originator will not get the NACK through n concatenated rings for n m-ps. For short data blocks this is virtually the same a

software solution 3. In general, concatenated networks will not all be rings, so the implicit solution is not generally suitable.

6. Conclusions

At present our inclinations are towards removing the header acknowledgement and combining it with the data acknowledgement. Whilst this results in low performance at present we expect performance to improve drastically with intelligent DMA interfaces. We are also exploring the Cambridge Protocols in the light of our own internetworking activities in the hope that some compromise might result in some form of ring standard.

Question Time

Questions centred around the re-invention of protocols in local area networks which were already well understood and implemented in wide area networks. Local networks differ in important ways from wide area networks, namely they have higher bandwidth, lower propagation delays and are more reliable. Thus protocol mechanisms such as windowing are overkill. Perhaps the best summary is to say the the same transport service interface is needed in both local and wide area networks, but that the mechanisms for achieving it will be different for different technologies.

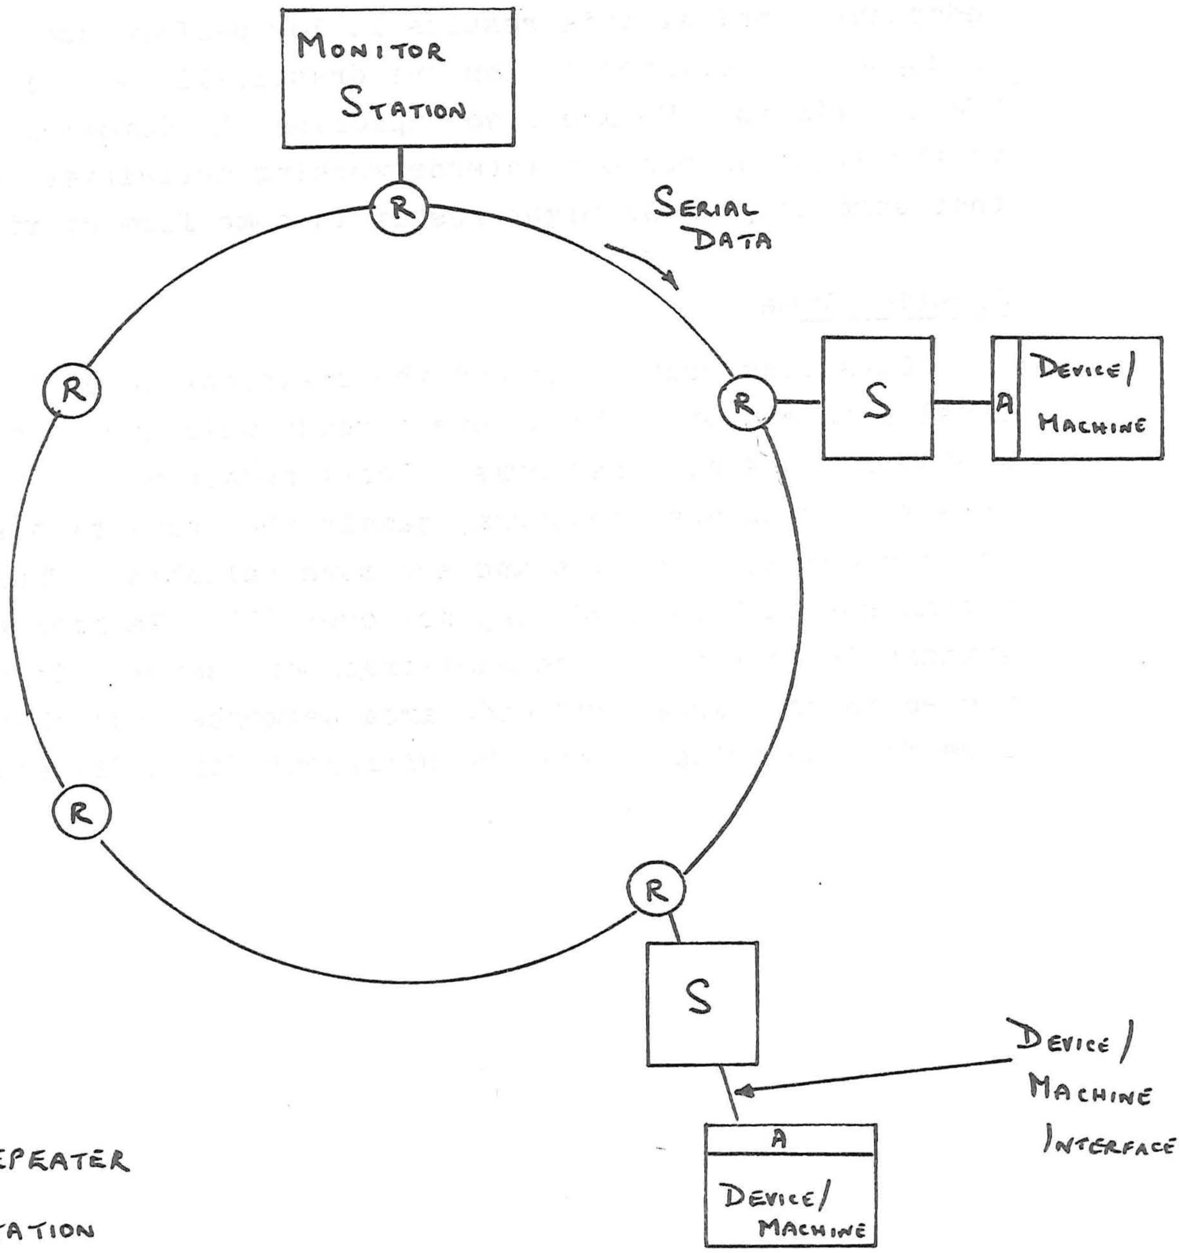

R - REPEATER
S - STATION
A - ACCESS LOGIC

FIGURE 1

STRUCTURE OF CAMBRIDGE RING

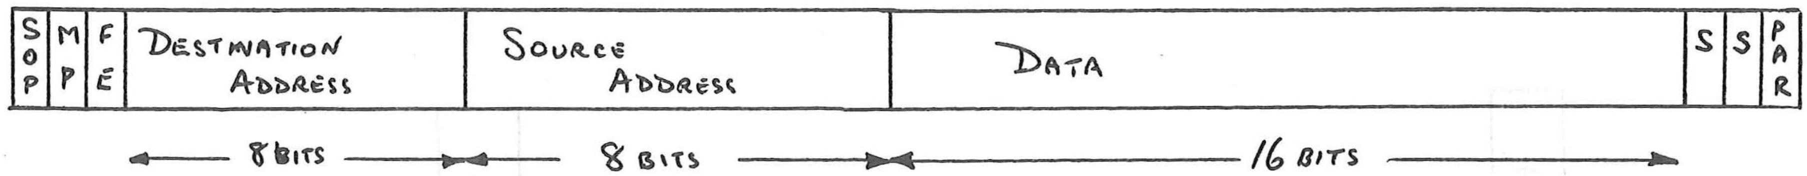

SOP - START OF PACKET
MP - MONITOR PASSED (MAINTENANCE)
FE - FULL/EMPTY
S - RECEIVED STATUS:
ACCEPTED
BUSY
NOT KNOWN
REFUSED

PAR - PARITY

FIGURE 2

MINI-PACKET STRUCTURE

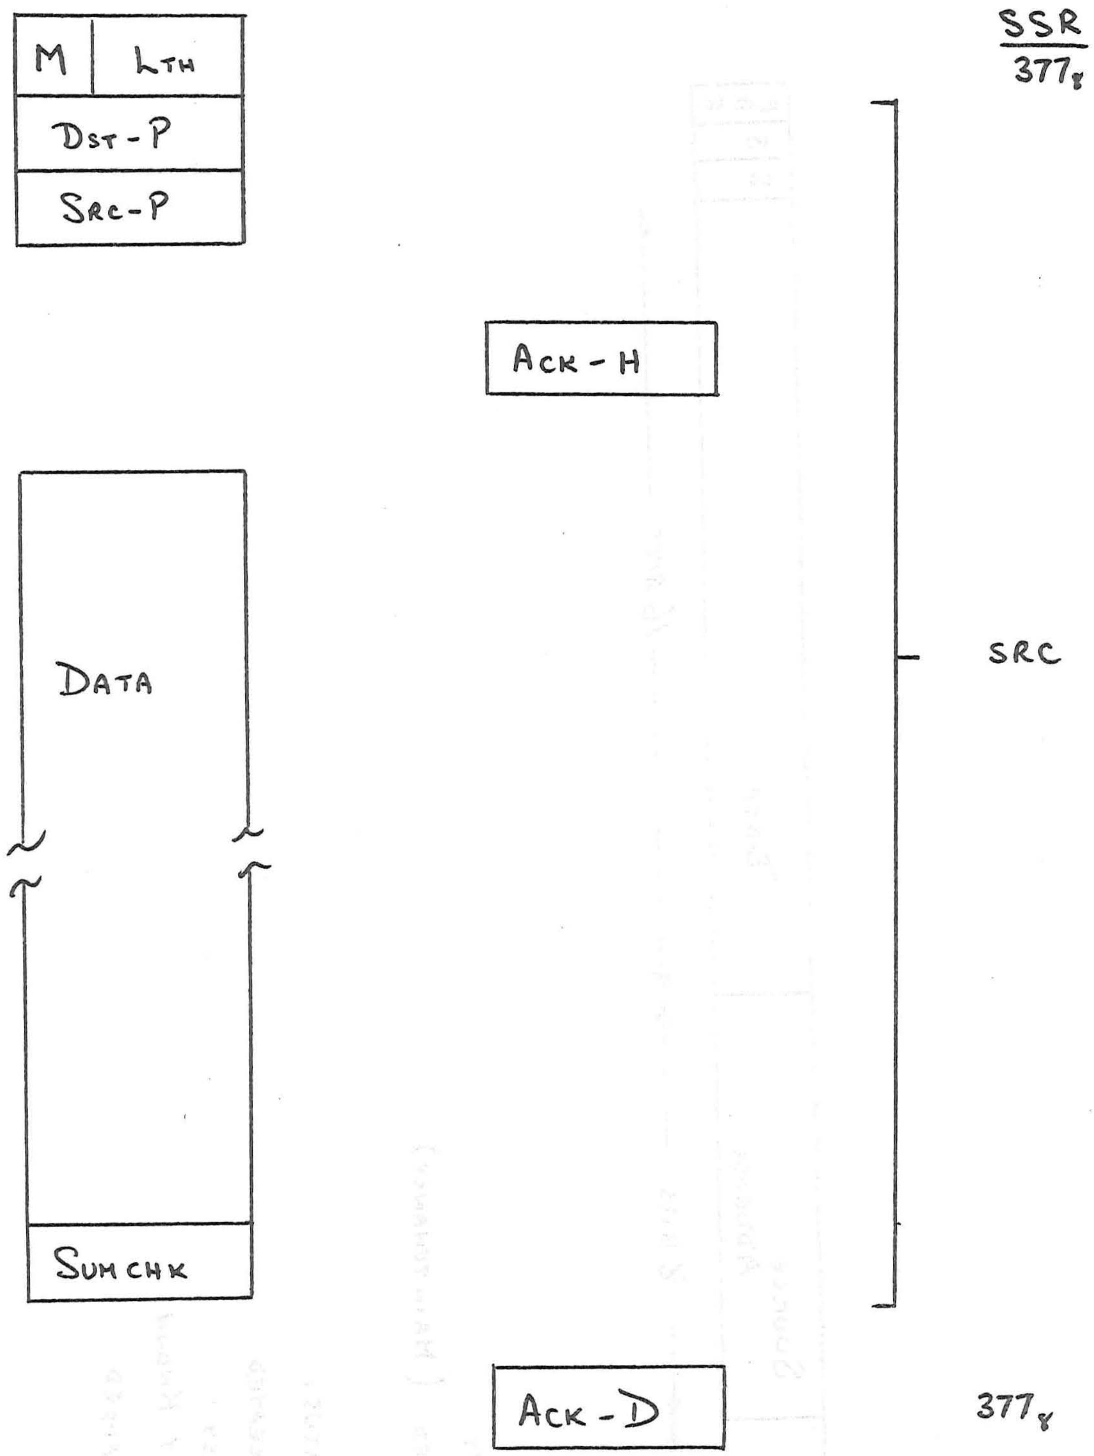

FIGURE 3

PROPOSED PROTOCOL

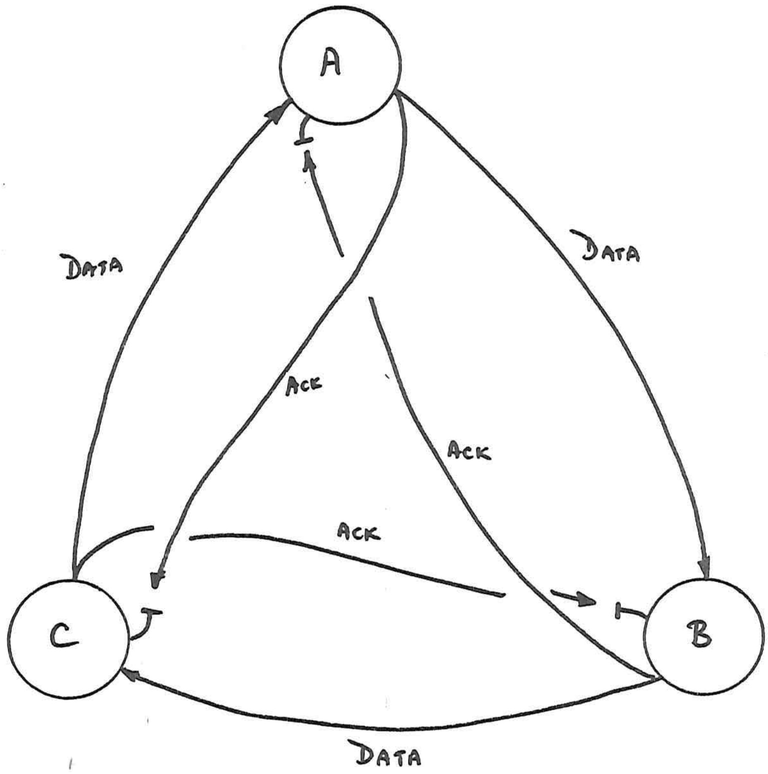

FIGURE 4

PROTOCOL DEADLOCK

INTERNETWORK GATEWAYS TO SUPPORT TRANSPORT SERVICE

ANDREW HINCHLEY

University College, London

INTERNETWORK GATEWAYS TO SUPPORT TRANSPORT SERVICE

by A.J. Hinchley

Department of Statistics and Computer Science
University College London

1. Introduction

The term "gateway" is now in common usage in networking to define a function where either the protocols, addressing or ownership of a communications path change. This paper is, however, restricted primarily to situations where potentially all three aspects may change at the gateway. Such gateways are part of the design indicated by the "yellow book" transport station definition.

The objective of constructing internetwork gateways of this type is, of course, to support an end-to-end transport service: - a multiplexed communication byte stream conveyed by any appropriate lower level network mechanisms, and by virtue of an extensible address scheme conveyed to any destination point.

2. International PTT Gateways

Before looking in detail at the general gateways just described, it is worth looking at a more limited gateway design: -enshrined in the CCITT specification X75 for the connection of X25 networks.

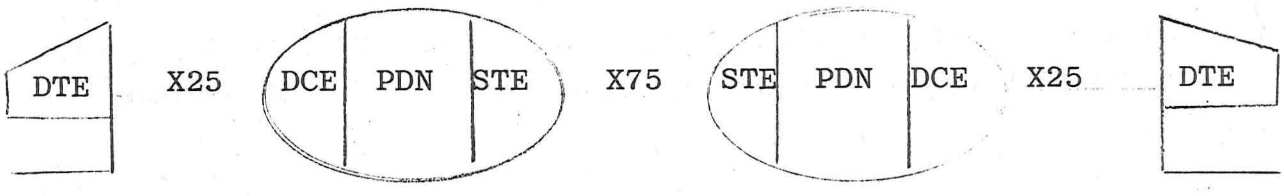

International Public Data Network System

In this design, each X25 virtual call requested by a user results in a corresponding call being set up in each of the X25 networks supporting the internetwork call. It is therefore the responsibility of each gateway to receive an indication of a new call request, and to set up an appropriate X25 call in the next network to either the final destination, or the most appropriate gateway. It is then a continuing responsibility of the gateway to support the mapped call on either side of the gateway and allow data to be transmitted on the duplex circuit. Signals such as X25 reset must also be recognised and propagated appropriately on the concatenated path.

It is worth noting the point that the X75 specification does not include the STE (Signal Terminating Equipment) shown on the diagram. The functions of the STE, and to an extent, some of the other gateway functions have to be inferred from X75.

The main results of the design are:

- The error rate of the concatenated call is the sum of the error rates of the component calls.
- The probability of failure of the call is the sum of failure probabilities of the component calls. (No provision exists to remake broken calls.)
- 128 byte packets only are supported (No provision for assembling or fragmenting packets at the gateway.)
- The requested throughput class is attempted to be satisfied by all component parts of the path.
- Flow control is by back pressure (No end-to-end significance.)

Additionally no reverse charging is supported in order that tariffs can be used which are different in each direction. IPSS already make substantial use of this.

The resulting service is therefore in theory accessable by use of an unchanged local network X25 DTE call without any changes in procedure subjects to the conditions mentioned above.

3. Gateways as Defined by the "Yellow Book" Transport Service

These gateways are comparable to the X75 gateways just described in that they support concatenated connections over an inter-network route, and the throughput delay and reliability of the connection is given by the properties of the individual connections which sum to provide the transport connection.

The gateways differ in that:

- Different network protocols can and will exist on different sides of the gateway.
- A byte connection is supported irrespective of packet sizes of contributing networks.
- End-to-end signals of the transport service are recognized by the gateway and generally both acted on and passed through unmodified.

Additionally, as the gateways may exist between any two networks, they will be built and controlled by varying organisations, unlike X75 gateways which are accessable only to PTTs. Additional access control features can thus be optionally built-in for those wishing to filter traffic through the gateway.

A final major difference is that whilst X75 gateways subscribe to a fixed limited CCITT addressing scheme (X121), "yellow book" gateways can support unstructured, flexible and extensible addressing convention. This is discussed in more detail later.

A likely UK internetwork picture might be:

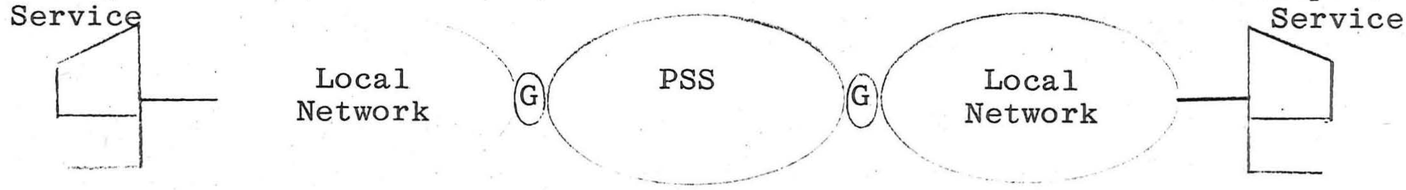

3.1 Gateway Functions

The functions of such a gateway are thus summarized:

- Interface appropriately to the two (or more) networks between which the gateway is supporting communication
- Set up and maintain connections subject to parameters such as throughput, etc. which may be specified
- Support the transport addressing scheme to route individual requests through the next network
- Recognize and act on the control signals of the transport service (The yellow book provides state tables for this.)
- Operate flow control across the gateway by "back pressure"
- Deal both with failures of connection attempts and failures of connections during data transfer

These functions map into a gateway picture as follows:

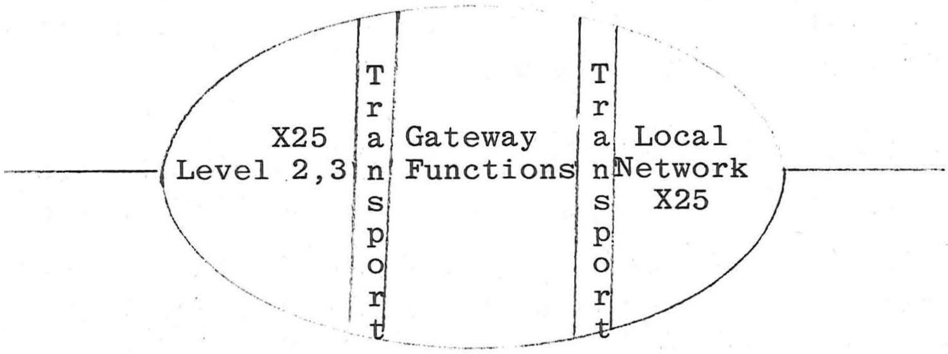

4. Application within the UK

The current Post Office regulations would require that, in general, traffic outside one organisation should travel by PSS or equivalent national network. The picture shown earlier, then, of two local networks exchanging traffic via PSS is the appropriate normal situation to comply with current legislation.

A general picture of cooperating users such as British universities would appear:

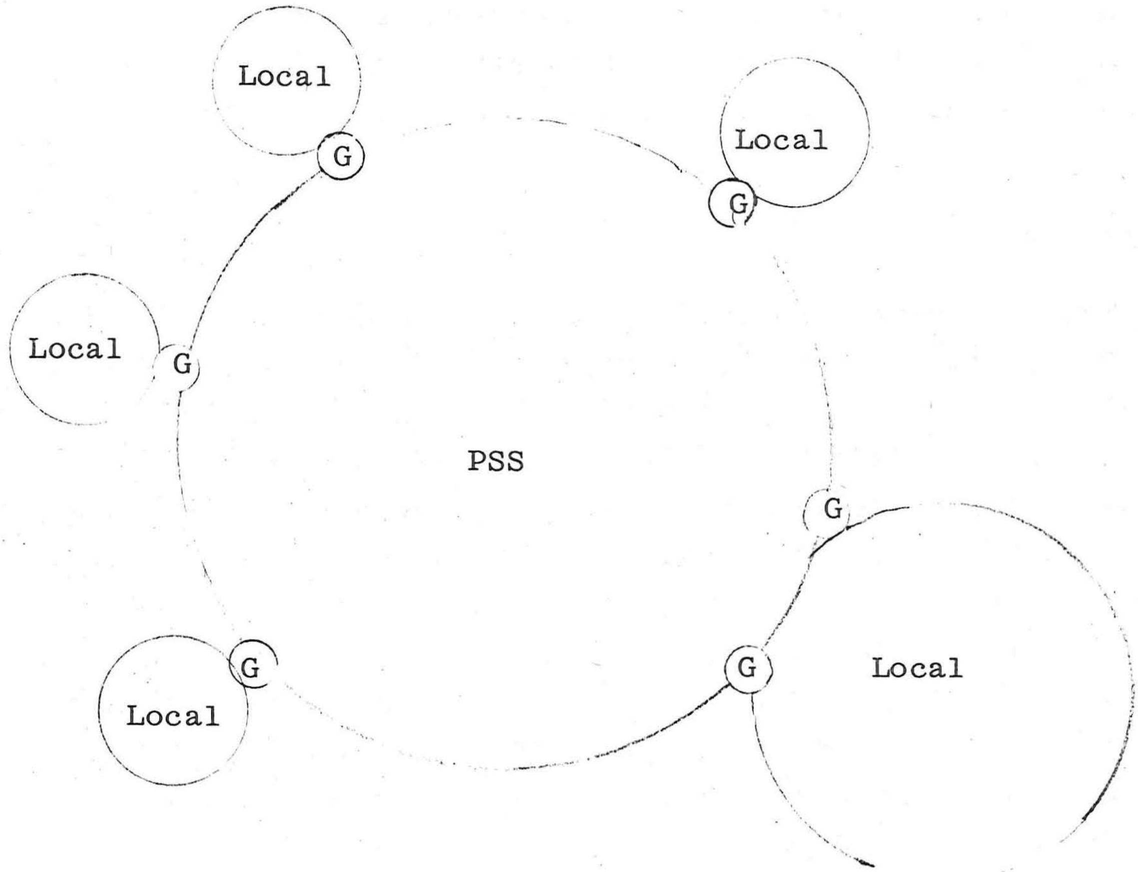

Two specific points can be noted. Firstly, regional switches as described at earlier networkshops can be essentially transparent in this picture if they operate within PSS addressing and protocol conventions. Secondly, within a region or local environment, a choice exists on how many gateways are used to connect that local network to PSS, and also internally if the local "network" is several physical networks, a choice also exists on whether to make the gateways that join those network "yellow book" gateways or to make them invisible at this level.

5. Some gateway functions in more detail

The functions broadly described in section 2 earlier are now examined in a little more detail to give some idea of the gateway implementation requirement.

5.1 Addressing

The transport service provides a very openended structure for addressing which is explained in detail in the yellow book, and was also the subject of a paper at the previous networkshop. The path to be taken over an inter-network route can be defined by the transport service user in terms of the gateways which lie between the networks, together, of course, with the name of the host/service which is required at the destination. If used in this way the gateway responsibility is then to deal only with the next network address of the pathname. In terms of the gateway picture shown earlier, the network specific parts of the gateway can handle onward calls into their network, leaving the common gateway function to be minimal.

5.2 Connection failures

In an environment where the source routing just described is operating, it is useful for the gateway which cannot achieve a successful connection to identify itself back to the user of the transport service. It may be possible to specify a different route to the destination to avoid the problem detected on the first route. Provision is made in the transport service for a gateway to give a full text description of the failure reasons. This cannot however be achieved in an X25 clear packet and thus can only apply if the component calls of the service are established to each gateway prior to confirmation being available on whether the connection can be supported.

During data transfer, an X25 reset will be propagated in both directions of the connection, while total loss of a component

call of the connection will result in disconnection since no procedures exist to re-establish the connection .

5.3 Flow control

The transport service as defined does not support end-to-end flow control and therefore, the sum of the available flow control mechanisms in each network through which the connection is established, operate to provide control over the entire connection. When no further data can be accepted it is up to the gateway to get this information to the sender by "back pressure". For this to happen effectively, the gateway should limit the buffer available for any single connection. In a congested situation, each gateway on the path will have a number of full data buffers for that connection, which is obviously wasteful of resources.

To maintain a requested throughput level on a particular connection, a gateway will however need to provide sufficient buffers to maintain the necessary receive transmit windows.

5.4 Access control

As a general rule it is unclear that access control is always required at this level. The higher levels above the transport service: - interactive, file transfer, job transfer, etc. provide protection for access to computational resources. There is, however, a specific requirement to protect communication resources where the cost is incurred elsewhere than the originating user. This applies, for instance, for out-going calls from a local network to PSS. Two approaches exist: - either additional user specific information must be included at connection initiation, interpretable by the gateway, or alternatively at connection closing, information on call cost should be cumulated. It can be returned to the system supporting the originating user which can direct the information to a suitable accounting program. The preference will depend on whether the objective is primarily to block unauthorized use, or to allocate incurred PSS costs across a user base.

6. Gateway implementation

Many of the features described will be common for any gateway, in particular the PSS support and all the central gateway functions. Differences may occur in the local net support and local net addressing. Optional access control features may also vary.

The gateway function can be combined with other communications oriented functions such as: local switch, PAD and protocol conversion. However, the placing of the gateway in a separate processor will lead to easier testing, be more reliable, and allow easier replication in other environments.

ACTIVITY WITHIN THE BSI

PETER LININGTON

Data Communication Protocols Unit

Activity within the BSI

P.F.Linington,
Data Communication Protocols Unit

The main thread of activity within the BSI committee DPS20 has been concerned with participation in the activities of ISO/TC97/SC16. The work of SC16 will be reviewed by the TC97 plenary meeting in November 1979.

At the plenary, SC16 will propose that the following projects be pursued:

Reference Model

Physical Layer Services and Protocols

Data Link Layer Services and Protocols

Network Layer Services and Protocols

Transport Layer Services and Protocols

Session Layer Services and Protocols

Virtual Terminal Protocols

File Transfer, Access and Management

Job Transfer and Manipulation Protocol

Management Protocols.

The direction of the BSI and ISO standards work in this area will depend critically on the response to these proposals.

REPORTS ON MACHINE RANGE NETWORKING ACTIVITIES

- PDP-11 (PAUL KUMMER, Daresbury Laboratory)

- DEC SYSTEM 10 (MIKE SAYERS, Hatfield Polytechnic)

- ICL 2900 (JOHN THOMAS, S.W.U.R.C.C.)

PDP11 NETWORK USERS MEETING - 3RD MEETING

UNIVERSITY OF LONDON COMPUTER CENTRE, 3 JULY 1979

Present:		
	P S Kummer	DL
	I Service	York
	I C Wand	York
	H C Kirkman	ULCC
	R J Ward	Bradford
	I R Dallas	Kent
	R Blake	Essex
	A C Peatfield	ULCC
	A J McLaren	Bristol
	A Hillman	NPL
	R Bismuth	Swansea
	C Bradbury	UCL
	R Lewis	Exeter
	M Guy	Cambridge
	R Rosner	JNT

Implementations

UNIX

York are still intending to do this and are waiting for completion of the contract arrangements with SRC. The project will take two years and will include all levels upto and including FTP. No decision has been taken yet on terminal protocols.

The Royal Institute of Technology, Stockholm, have done a lot of work on UNIX/X25 using an LSI11 front-end processor and a 16-bit parallel interface to the PDP11. York will be getting documentation at the end of the summer.

It appears that inboard networking software on an 11/45 (driving its synchronous port on program interrupts) uses 60% + of the CPU.

Version 7 of UNIX is due soon - this is a major rewrite.

RSX11-M

Bristol - still interested but no hardware.

UCL - have done RSX11-S upto level 3 but extension to use memory management is a lot of work. This implementation could run on an outboard LSI11 and use a 16-bit parallel interface to the PDP.

There was a short discussion on the cost of an outboard LSI11 solution (which may also have line speed limitations), and production and maintenance problems. It was agreed that there should be an approved list of parts which should be easily maintainable. The JNT would supply the software for this list of parts.

If an outboard LSI11 solution using the UCL software is agreed, then software will be needed to run under RSX11-M and drive the 16-bit parallel interface. The format of the data on the 16-bit link is defined. This may be a contender for a commercial contract.

Swansea - writing FTP (IFC subset) for RSX11-M and RSX11-D. Completion is scheduled for the end of the summer. Written in MACRO. There may be a problem with machine time. This implementation could be used as a starting point for a fuller implementation.

RT11

A proposal was received from Exeter and Bristol. There may be hardware problems although the manpower is probably available. A copy of the proposal is attached to this summary.

RSTS

Hatfield are no longer interested in this and as nobody at the meeting expressed a requirement for RSTS/X25 further effort has been deferred.

PADS

ULCC are producing a provisional PAD specification as a starting point. This will be circulated to interested people and a meeting organised (contact A C Peatfield at ULCC for details). It was noted that a working group of SG3 has produced a document to define what a PAD should do.

A standalone PAD will probably be a non-DEC OS but it should be possible to drop the software into a DEC OS.

Transport Service

A final draft (?) is expected by the end of August. It was agreed that all implementations should provide TS and X25 level 3 interfaces.

Peter Higginson at UCL is heading a group working on the compatability between X29 and TS.

FTP

This should be to the subset defined by the FTP implementors group.

JTP

Work on this is progressing slowly and is still at the function level.

Steering Groups

It was agreed that small steering groups should be set up to monitor the progress of the various projects. Membership was agreed as follows:

RSX11 - JNT
DL

RT11 - JNT
ULCC
Exeter
Bristol

Any comments on the RT11 proposal should be sent to the Exeter and Bristol people as soon as possible.

DEC's Position

An announcement from DEC about PSS compatible implementations of X25 is expected in September.

Next Meeting

The next meeting has been arranged for Friday, 19 October 1979 at ULCC. Please note the change of date from the proposal at the meeting.

P S Kummer

24 July 1979

DECsystem-10 Networking Status Report

DECsystem-10 networking projects are co-ordinated by a group representing the interests of Computer Board and SRC DECsystem-10 sites which meets under the aegis of the JNT. A number of projects have been identified as necessary and most have now been specified and development projects have been funded. A summary of the current state is given below.

1. ANF10 to X25 Gateway

   First stage, connected to SRC net and supporting ITP is now working. The next stage is to add

   a) PSS compatibility

   b) Triple X interactive terminal protocols.

   It is expected that stage two will be ready for testing by the time PSS service is available.

2. FTP for DECsystem-10

   A functional specification, describing user and operator interfaces and FTP options required has been prepared by CAP after extensive discussions with members of the DECsystem-10 networking group. A proposal to implement this specification has been received from the University of York and support is to be arranged through the JNT.

MDS/VRAC
24.10.79.

2900 RANGE - PROGRESS REPORT

Significant progress has been made towards setting up a joint SWURCC/ ICL project. With the full support of the JNT this project will ensure that objectives specified to manufacturers by the Computer Board for the supply of networking software on new mainframes will be met.

During the last three months of this year SWURCC technical staff will be working alongside relevant ICL staff and towards Christmas it is expected that well defined plans and full details can be made publicly available.

CLOSING REMARKS

ROLAND ROSNER

Joint Network Team

CLOSING REMARKS

This workshop has been marked by more active audience participation in discussions and by a trend, on the part of some presenters, to describe problems rather than to disguise them. Both developments are to be encouraged, for workshops are not conferences at which papers are traditionally confined to successes presented without the warts. Hopefully, we can learn from each other's mistakes and it may even be worth soliciting future contributions which specifically aim to describe failures and experiences with incorrect techniques.

A cornerstone of network development is the emphasis placed on standards and the workshops have offered detailed coverage both of the definitions and of implementation strategies. However, there seems to be a growing danger of paying no more than lip-service to the standards by tampering with them to match local circumstances. Phrases such as "subsets", "supersets", "X25-like", "most of the FTP" have been used with alarming frequency during this workshop. Perhaps no further warning is needed on the problems which could be caused by too loose an interpretation of standards. This raises the whole issue of testing for conformity and both the JNT and the DCPU will be concerned with the provision of appropriate facilities.

There has been much discussion on the need to conserve scarce manpower by trying to pick up components and software developed elsewhere. In particular, it has been suggested that, for multiple operating systems running on the same machine range, protocol implementations might share a significant amount of common code. This has been extended to encompass the proposal that high level languages might be used to create protocol implementations which would be portable across machine ranges. The main debating points are which language to use and how much of a protocol implementation is operating system-dependent. The JNT will initiate activity in this area.

The production of components and packages is taking place largely by means of contracted projects and the problem of monitoring progress and defining what constitutes completion of the task has been raised. The demand for robust products dictates that a more detailed elaboration of a project's objectives may be required than has hitherto been the case. There will also be a need for the setting of project milestones (with timescales) and an acceptance procedure at the conclusion of the project. The JNT will lay down some general guidelines to satisfy these criteria.

Some hands have been raised in horror at the magnitude of the task which faces the community in moving from current arrangements to the brave new networking world. Without attempting to minimise the scale of what is being attempted, it is possible to detect some reluctance to devote time to the drafting of plans which define not only the end objective but also the transition phase. It is unrealistic to expect things to turn out right on the night or to await plans to be handed down from on high. Centres must themselves contribute to the formulation of these plans with the JNT acting as a catalyst and ensuring that the needs of all the funding bodies are taken into account.

Dr R A Rosner
Joint Network Team

22 October 1979

APPENDIX 1

PROPOSED PILOT INSTALLATION OF A CAMPUS NETWORK

HOWARD DAVIES

South West Universities Computer Network

(Received too late for inclusion in the proceedings of Networkshop 4)

SOUTH WEST UNIVERSITIES COMPUTER NETWORK

NETWORKSHOP 4

Issued by
H E Davies

Proposed Pilot Installation of a Campus Network

14 August 1979

---

1. Introduction and Summary

In its Capital Expenditure Estimates submitted to the Computer Board in December 1977, the South West indicated a requirement for campus switches to be installed progressively at each site in the region.

Consideration of the current level of development of manufacturers products suggests that it would be imprudent to rely on a commercial product being fully integrated and proven at that time. It will be necessary, from a national viewpoint, for at least one pilot implementation to be carried out before user service is committed to such a network. A project for a pilot installation would be required to integrate the system in a realistic mix of connected machines and terminals, to check out the total system, and possibly carry out a residual development to rectify defects and omissions. This paper proposes such a project. It would involve the selection, purchase and installation of a campus packet switching exchange (CPSE) and packet assembler/disassembler (PAD), the interconnection of a representative set of existing site and departmental computers and a demonstration that the whole communications system met its specification in terms of traffic capacity, reliability and operability before the first production service is required in summer 1980. The pilot installation would be in Exeter, the site which is offering to provide half the necessary manpower. Staff at other South West sites will be fully occupied during the next two years in selecting new mainframes and working to establish user services, including networking services, on them.

As a demonstration of a particular form of campus networking system, the project is seen as being of national benefit. It would also exploit the expertise in the use of X25 and other national networking standards which has already been built up in the NIP development project and in connecting the MULTICS system. Assuming that the pilot installation is successful, we would expect similar equipment to be installed progressively at the other South West sites. There is an urgent need for a production system to be installed in Bristol early in 1980 and for a further installation at Bath soon afterwards. Detailed cost estimates will not be available until formal offers are obtained from potential suppliers, but current indications are that £40 - £50,000 will be required for the first installation. The manpower required amounts to 4 man years and will be provided from existing establishments.

2. Objectives

The primary objective of the project is to install and commission a pilot campus network based on a central switching device so that from the middle of 1980, sites in the South West which have the need can employ a campus network as a routine part of their production service. Work is already under way on the development and installation of communications software based on national and international standards, both on mainframes and on minicomputers of different types which must be able to communicate with them. The proposed project is therefore mainly concerned with the CPSE and PAD and more specifically includes:

a) the selection of suitable equipment,
b) the validation of the manfuacturer's newly released software (though a large part of this should be done by the Post Office as part of its procedure for accepting a connection to its own services),
c) demonstrating that the specified throughput can be achieved,
d) ensuring that sufficient control and measurement software is available to provide efficient management of site communications,
e) producing a report describing the experience gained, the difficulties encountered and the means used to overcome them.

An additional benefit will come from proving the communications software in the connected computers.

All the hardware and software needed for the project either exists already, will be bought "off the shelf" from a supplier or is being provided as part of other existing projects. In principle, all these components should fit together to provide a working system without any further development, but experience shows that in an area as technically complex as this one, the final assembly and checkout of the system requires conscious and determined effort if smooth operation is to be achieved by a pre-defined deadline.

It is our intention to work with the Computer Board's Joint Network Team at all stages of the project so that national standards are respected, so that the work done is of national benefit and so that the South West benefits from related work being carried out elsewhere.

3. The Case for Campus Networks

3.1 The General Case

During the last decade, large numbers of computers have been bought for such jobs as the control of experiments, data acquisition, local editing, small computational jobs and local teaching. Besides their low cost, small machines have an advantage over mainframes of being convenient to use. They usually have a simpler user interface, they are continuously available, and their users have complete control over their operation.

Simple microprocessor systems are now so cheap that it is surprising that relatively few are in use. The conventional explanations are that software and peripherals are still required and the cost of these is the same as for classical machines. This represents only part of the story.

Several examples of linked systems demonstrate that mini- and micro-computer systems can be most effectively exploited by connecting them to mainframes and making use of their power and cost effectiveness in supporting facilities such as large file stores, cross-compilers, debugging aids and special software packages. Although suitable hardware and software techniques are now available for providing generalised communications systems, there are still two obstacles which have inhibited their growth.

The first is the lack of communications standards which has prevented software producers from risking a large investment in communications software on one system when there is no guarantee that communication with other systems would be possible. The second is the continuing desire of manufacturers to supply their customers with complete systems and their lack of interest in sharing an installation with their competitors.

The first of these impediments will be removed by the introduction of the Post Office's PSS which will act as a catalyst for the development of standard interfacing softare on almost all machine ranges. Manufacturers' attitudes will take longer to change but pressure is already being put on them to provide standard communications interfaces within their operating systems in the same way that they already offer standard hardware interfaces to terminals.

Once this happens, the effects will be dramatic. Staff of computer centres will no longer be constrained to choose the single manufacturer whose equipment provides the best match to their users' needs. Instead, they will be able to install sets of more specialised systems, perhaps collaborating with other centres (as already happens in the South West) to share capacity and still take advantages of economies of scale. Individual users will benefit even more. They will be able to select a personal computer which suits their taste, choose from a much wider range of mainframe services than is currently possible and combine the use of local and remote facilities in the way that suits them best.

This transfer of choice of service to the user represents the next major step forward in the development of computing technology. Powerful data communications systems acting at national and local levels are essential if the advance is to be made. The national system will be introduced this year but will have no value unless local services (campus networks in the case of Universities) are also available to provide access points for individual users.

3.2 Requirements in the Region

Even within the present System 4 network, local communications play a vital role in giving users access to their local mainframe and to other machines in the Network from terminals and RJE computers dispersed around the sites.

As the demand from users for remote access to campus and regional computing facilities increases, a large proportion of the data traffic within each site will continue to be in and out of its local mainframe with a smaller fraction of terminal and RJE traffic going to mainframes at other sites and national centres.

Although the number of connections of other kinds, for example for file transfer between two minicomputers on the same or different sites, is expected to be very small as a proportion of the total, a generalised communications system, potentially allowing connections between any pair of computers on the same or different sites is required.

In the present system the mainframes themselves handle all switching of data to or from the computers at their site, but nowadays this function can be handled much more cheaply, efficiently and reliably by a mini-computer or special-purpose device. New mainframes will be connected to the campus network but will not be directly involved in its operation or control.

In the South West, the timetable for the installation of campus networks is therefore governed initially by the mainframe replacement program. As each System 4 computer is closed down, the local communications service it provides must be taken over by a new system. Otherwise, equipment installed in departments to work as RJE stations or terminal concentrators will lie idle, and there will be great pressure to take temporary, ad hoc (and expensive) short term measures to maintain services that are now regarded as essential.

The first System 4 replacement machine has already been installed and the South West needs to introduce at least one full campus network into service by the summer of 1980. The introduction and gradual buildup of Multics services pose two specific communications problems:

- connection of RJEs and remote I/O service
- resolution of contention for terminal ports

Clearly, ad hoc solutions to these may be achieved in the requisite timescales; RJEs might be attached by 2780 emulator (expensive), terminal simulation (restrictive), or by BOSS implementation on Multics (retrogressive). The 3:2 ratio of terminals to physical ports (envisaged by late 1979) might be handled by a Gandalf type device, or by the Daresbury-like concentrator (diluter system) or by a multiplexor with demultiplexing software in Multics. None of these solutions would make a significant contribution to the Universities' real requirements; rather they would deprive them of manpower and funds which could be better employed.

The interconnections between computer systems foreseen in a few years' time are shown in figure 1. All the sites in the South West have requirements which follow the same general pattern but which differ in detail.

4. Form of Campus Networks Proposed

The South West has followed the initiative of the Network Unit in investigating three forms of local area Network; Ethernet, the Cambridge ring and a central switch as proposed by several manufacturers in response to the Network Unit's specification. Ethernet is now believed by both its developers (QMC) and sponsors (Rutherford Laboratories) to be inferior to the Cambridge ring, and has therefore been discounted The ring as implemented at Cambridge and Canterbury is not a manufacturers supported product and is, furthermore, considered to be more appropriate to the needs of a compact campus or a large department rather than to a large extended site. They are therefore not thought to meet the South West's requirements at present. The CPSE approach on the other hand is the one which is the nearest to giving a proven commercially available product.

The CPSE would provide an X25 switching function and so all computers connected to it must offer X25 protocol support. In particular cases, for example if an X25 interface has to be produced for an existing machine, it may simplify the development if special facilities of X25 such as permanent virtual circuits are used. The CPSE must therefore be able to handle all aspects of the X25 protocol. The CPSE should ultimately be connected to the Regional packet switch exchange (RPSE) thus giving access to the Post Office's PSS. A direct connection between the first CPSE and PSS would be set up in the meantime.

Other computers to be linked to the CPSE at any site would include:

a) one or more host computers including the site's mainframe,

b) at least one departmental mini-computer supporting some combination of terminal connection, file transfer and RJE to the host computers and to Regional and National services.

Facilities for terminal concentration and packet assembly/disassembly (PAD) offering at least X3/X28/X29 support for character mode terminals must also be provided either within the CPSE or as a stand-alone system.

5. Costs

Between July and December 1978 each South West site produced estimates of the traffic to be handled by a campus PSE at peak periods in about four years time. The summary estimates for Bristol and Exeter are given in appendix 1.

These estimates were sent to a number of manufacturers who were believed to be capable of offering off the shelf equipment suitable for a campus PSE. Discussions were held with four manufacturers, Systems Reliability Ltd, GEC Computers, Plessey and Logica. The Post Office was also approached indirectly via the Network Unit.

Briefly, Systems Reliability turned out not to have suitable equipment. Plessey are able to deliver a system on the planned timescale but the cost would seem to be the highest; their equipment is oriented towards the needs of very large scale Networks and would be expensive on the smaller scale of a single campus mainly because of the need for a Network control centre costing about £100,000 at one of the sites. Both GEC Computers and Logica seem capable of supplying the necessary equipment at prices comparable with the budgetary figures that have already been presented to the Computer Board.

The Post Office has expressed an interest in supplying equipment for the campus switch. Their policy on supporting this kind of service seems however not to have been defined yet and it is unlikely that they would meet our time scale. They would however be consulted again at the formal tender stage.

On the basis of the information already supplied by GEC and Logica, the capital cost of the new equipment required should be about £40 - £50,000 for the pilot installation.

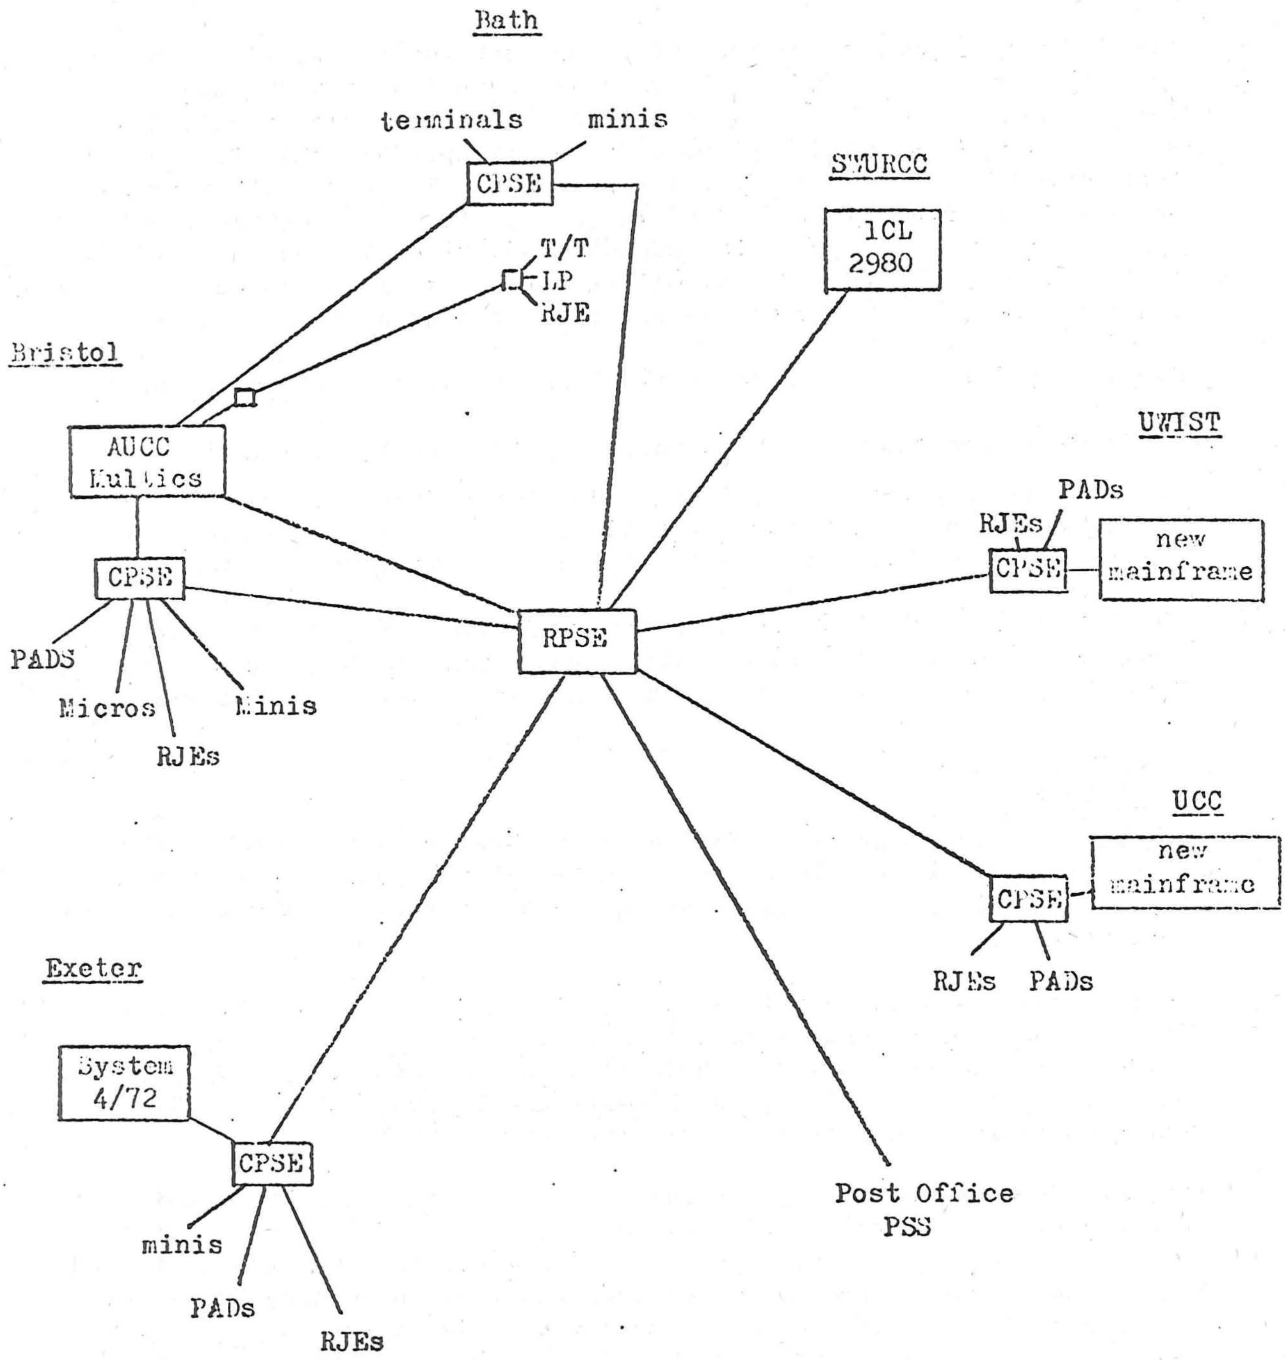

Fig. 1: Future Connections between Computers in the South West

Appendix 1: Estimates of Data Transmission Requirements at South West Sites

1. University of Exeter

A. Lines

1. Number of lines at each data rate 3 x 48 kbaud; 8 x 4.8 kbaud

2. Number of lines at each protocol 8 x X25; 3 x local(BOSS)

B. Throughput

1. Gross data throughput of the switch* 120 k bit/sec

2. Actual data flow on each line as a percentage of nominal data rate*
   - input to switch
   - output to switch

   48 kbaud : 32%, 26%, 26%
   4.8 kbaud: all 50%

3. Packet size

- average	625	or	435
- maximum	1024		512
4. Packets per sec*	192		276

5. Number of call establishments per sec* 0.27

6. Number of concurrent calls in progress* 85

C. Reliability

1. Availability of total switch >99.6%(½ hr loss per 3 shift week)

2. Availability of system to a given user >99% (1.2 hr loss per 3 shift week)

* During the peak ½ hour

A. Lines.

1. Number of lines at each data rate

1 or 2 x 48 kbaud

1 x 9.6 kbaud

38 x 4.8 kbaud

2. Number of lines at each protocol 40 or 41 at X25.

B. Throughput.

1. Gross data throughput of the switch* 158 kbits/sec

2. Gross data flow on each line as a percentage of normal rate*

- Input to switch

1 48 kbaud ∓ 98% or 2 48 kbaud:49%

1 9.6 kbaud = 65%

10 4.8 kbaud ∓ 5%

6 4.8 kbaud = 15%

12 4.8 kbaud = 25%

10 4.8 kbaud = 10%

- output from switch

1 48 kbaud ∓ 46% or 2 48 kbaud 23%

1 9.6 kbaud ∓ 34%

10 4.8 kbaud ∓ 45%

6 4.8 kbaud = 50%

12 4.8 kbaud = 25%

10 4.8 kbaud = 10%

3. Packet size

- average	356		294
- maximum	1024	or	512

4. Packets per sec* 440 532

5. Number of calls established per sec* 0.79

6. Number of concurrent calls in progress* 199

* During the peak ½ hour.

C. Reliability.

1. Availability of total switch	99.8%	(¼hour loss per full week)
2. Availability of system to a given user.	99.7%	(½hour loss for full week)

APPENDIX 2

NETWORKSHOP 5 PROGRAMME

Networkshop 5 - Programme

Wednesday 19th. September

Chairman: Ian Dallas (Kent)

09.00	Introductory remarks - Brian Spratt (Kent)
09.15	Report from the Joint Network Team - Roland Rosner (JNT)
09.40	Report from the Data Communication Protocols Unit - Keith Bartlett (DCPU)
09.45	Progress on PSS - Pat Morrison (Post Office)
10.15	Discussion
10.30	Coffee

Chairman: Brian Spratt (Kent)

11.00	X25 Level 3 P.O. Technical Guide - Mike Sands (Post Office)
12.00	Testing X25 Implementations - John Horton (GEC) and John Thomas (SWURCC)
12.30	Lunch

Chairman: Ken Heard (AERE Harwell)

13.45	Final Draft of the Transport Service - Peter Linington (DCPU)
14.30	How to use old kit in a modern networking environment - Tony Peatfield (ULCC)
15.15	Tea

Chairman: Morley Sage (Southampton)

15.45	Testing High Level Protocol Implementations - Alun Jones (CADC)
16.15	Progress on Network Projects -- South West - John Thomas (SWURCC) -- North West - John Rice (Liverpool) -- Midlands - Phil Harrison (Nottingham) -- RCO - Bill Hay (ERCC) -- SRC - John Burren (Rutherford Laboratory)
17.30	End of formal session
17.45	(Communications Access to IBM - The Cambridge Solution - Chris Cheney & Mike Guy (Cambridge) (ICL 1900 Access to X25 - John Salter (ICL - Dataskil)
19.00	Dinner
20.00	Honeywell Network Users Group Meeting
20.30	Prime Network Users Group Meeting

Networkshop 5 - Programme (Continued)

Thursday 20th. September

Chairman: Chris Cooper (Rutherford Laboratory)

09.00	Preliminary Report on Job Transfer Protocol - Mike Guy (Cambridge)
10.00	Triple-X and the Transport Service - Peter Higginson (UCL)
10.30	Coffee

Chairman: Chris Morris (Bristol)

11.00	Terminal Concentrators, Switches and PADs - Harald Kirkman (ULCC)
11.30	Progress on the Ring at Kent
	-- Hardware - Matt Lee (Kent)
	-- Putting the Ring into Service - Steve Binns (Kent)
12.15	Trends in Communications Hardware - Ray Chisholm (ERCC)
12.45	DEC communications products - Ian Service (York)
13.00	Lunch

Chairman: John Prentice (Loughborough)

14.15	The File Transfer Protocol - Problems and Progress - Dave Rayner (NPL)
14.45	Experience in the implementation and usage of the FTP - Paul Bryant (Rutherford Laboratory)
15.15	Tea (End of Formal Session)
15.45	Problems of Parallel Running - Jeremy Brandon (QMC) & Tony Peatfield (ULCC)
16.30	An Introduction to Triple-X - Peter Higginson (UCL)
19.00	Sherry reception
19.30	Conference dinner

Networkshop 5 - Programme (Continued)

Friday 21st. September

Chairman: Mike Wells (Leeds)

Ø9.15 The New NPL Network - Keith Bartlett (DCPU)

Ø9.4Ø A Mainframe-to-Ethernet connection - Peter Girard (Rutherford Laboratory)

1Ø.ØØ Protocols for Local Area Networks - Steve Wilbur (UCL)

1Ø.3Ø Coffee

Chairman: Tony Young (Durham)

11.ØØ Gateways - Andrew Hinchley (UCL)

11.3Ø DCPU and BSI activities on Protocol Standards - Peter Linington and Keith Bartlett (DCPU)

12.ØØ Reports on machine range networking activities

-- PDP-11 - Paul Kummer (Daresbury Laboratory)

-- ICL 19ØØ - John Salter (ICL - Dataskil)

-- DEC System 1Ø - Mike Sayers (Hatfield Polytechnic) et al

-- ICL 29ØØ - John Thomas (SWURCC)

13.ØØ Lunch

Chairman: Mervyn Williams (Network Unit)

14.15 General Discussion and Summary

15.ØØ Tea (End of Proceedings)

APPENDIX 3

DELEGATES

NETWORKSHOP 5 - DELEGATES

Ashton, Mr I.	Bradford University
Aspden, Mr J. P.	Newcastle University (NUMAC)
Aswani, Mr V. G.	University of Manchester Regional Computer Centre
Austen, Mr A.	North London Polytechnic
Barkataki, Dr P.	Teesside Polytechnic
Barney, Dr G. C.	Control Systems Centre, U.M.I.S.T.
Bartlett, Mr K. A.	Data Communication Protocols Unit
Binns, Mr S. E.	University of Kent
Blake, Dr R. G.	University of Essex
Bonney, Mr N.	Brunel University
Boodson, Mr A.	Leeds Polytechnic
Bradshaw, Mr R. G.	South West Universities Regional Computer Centre
Brady, Mr P. R.	University College, Swansea
Brandon, Mr J. P.	Queen Mary College, London
Bryant, Mr P. E.	Rutherford Laboratory
Burren, Mr J.	Rutherford Laboratory
Buttle, Dr A.	University of Exeter
Bye, Mr C.	University of Dundee
Caul, Mr D. J.	University of Kent
Charles, Dr B. J.	Joint Network Team
Cheney, Mr C. J.	University of Cambridge
Chisholm, Mr R.	Edinburgh Regional Computer Centre
Clark, Mr T. B. G.	University of Warwick
Clelland, Mr S. R. M.	Heriot-Watt University
Cooper, Dr C.	Rutherford Laboratory
Couzens, Mr J.	Imperial College, London
Dainton, Mrs D. M.	University of Glasgow
Dallas, Mr I. N.	University of Kent
Dand, Mr A. K.	South West Universities Regional Computer Centre

Davies, Dr H. E.	South West Universities Computer Network
Diediw, Mr A.	Bradford University
Drabble, Mr A.	University of London Computer Centre
Driscoll, Mr R. J.	Hatfield Polytechnic
Ellison, Mr H. D.	University of Salford
Findon, Mr P.	University of Aston in Birmingham
Garside, Mr M. J.	University of Kent
Gillman, Mr R. M.	New University of Ulster
Girard, Mr P.	Rutherford Laboratory
Glass, Mr B.	Digital Equipment Corporation
Greenwood, Mr F. T.	Newcastle University (NUMAC)
Gurteen, Mr	Prime Computers Inc.
Guy, Mr M. J. T.	University of Cambridge
Harrison, Mr P. S.	Nottingham University
Hay, Dr W. D.	Edinburgh Regional Computer Centre
Heard, Dr K. S.	AERE Harwell
Higginson, Mr P. L.	University College, London
Hinchley, Mr A.	University College, London
Holt, Mr A. D.	City University
Horton, Dr J. R.	GEC Computers Ltd.
Hughes, Mr D. J.	Institute of Oceanographic Sciences
Jack, Mr A.	Nottingham University
Jacob, Mr G.	North London Polytechnic
Jamieson, Mr J. B.	Strathclyde University
Jennings, Dr D.	U.W.I.S.T.
Johnson, Mr D.	University of Manchester Regional Computer Centre
Johnson, Mr M. G.	Westfield College, London
Jones, Mr A.	Computer Aided Design Centre
Jones, Dr P. S.	University of Durham
Kennington, Mr C. J.	University College, London

Kirkman, Mr H. C.	University of London Computer Centre
Kummer, Dr P. S.	Daresbury Laboratory
Larmouth, Dr J.	Salford University
Lee, Mr M. N. A.	University of Kent
Linington, Dr P. F.	Data Communication Protocols Unit
Linn, Dr J. A.	University of Aberdeen
Litchfield, Mr G. W.	Oxford University
Lynch, Mr N. P. C.	University of Reading
Mason, Mr P.	Sheffield University
Morris, Mr C.	University of Bristol
Morris, Mr J. E.	University of Leicester
Morrison, Mr P.	Post Office Telecommunications
Nokes, Mr S.	Honeywell Information Systems Ltd.
O'Sullivan, Mr D.	Queen's University of Belfast
Patel, Mr A. M.	Trinity College, Dublin
Peatfield, Mr A. C.	University of London Computer Centre
Phillips, Mr A.	University of Lancaster
Piotrowicz, Mr M.	University of St. Andrews
Powell, Mr C. J.	Avon Universities Computer Centre
Powell, Mr R. G.	University of Leeds
Prentice, Mr J. A.	Loughborough University
Pugh, Mr C. G.	Imperial College, London
Rayner, Dr D.	National Physical Laboratory
Rice, Dr J. D.	University of Liverpool
Richmond, Mr I. M.	Rothamsted Experimental Station
Roberts, Ms S.	University of Bristol
Rosner, Dr R. A.	Joint Network Team
Sage, Mr M. W.	University of Southampton
Salter, Mr J.	ICL - Dataskil Ltd.
Sands, Mr M.	Post Office Telecommunications
Sayers, Dr M. D.	Hatfield Polytechnic

Scolley, Mr A. J.	University of Stirling
Service, Mr J. D.	University of York
Sexton, Mr J.	ICL - Dataskil Ltd.
Spratt, Dr E. B.	University of Kent
Spridgeon, Mr D.	University of Hull
Stanier, Dr A. M.	University of Essex
Swindells, Mr W.	Interactive Computing Facility, U.M.I.S.T.
Thomas, Mr J. S.	South West Universities Regional Computer Centre
Vivian, Mr R.	University of Birmingham
Wallace, Mr T. L.	Prime Computers Research Ltd.
Wanless, Dr D.	University of Keele
Wanless, Mr S.	South Bank Polytechnic
Watts, Mr P.	South Bank Polytechnic
Wells, Professor M.	University of Leeds
Wilbur, Mr S. R.	University College, London
Williams, Mr A. H.	University of Bristol
Williams, Mr M. B.	Network Unit
Wood, Mr B.	CAP (London) Ltd.
Young, Mr A. A.	University of Durham

ATTENDING AS NON-DELEGATES

Kennington, Mrs E. J.	Attending with Mr C. J. Kennington
Sage, Miss	Attending with Mr M. W. Sage
Sage, Mrs M. W.	Attending with Mr M. W. Sage
Williams, Mrs S. D.	Attending with Mr M. B. Williams

APPENDIX 4

INSTITUTIONS

University of Aberdeen Computing Centre 150 Don Street Aberdeen AB2 1SQ	Dr J. A. Linn
University of Aston in Birmingham Computer Centre 15 Coleshill Street Gosta Green Birmingham B4 7PA	Mr P. Findon
A.E.R.E. Harwell Didcot Oxon OX11 0RA	Dr K. S. Heard
Avon Universities Computer Centre University of Bristol University Walk Bristol BS8 1TW	Mr C. J. Powell
University of Bradford Computer Unit Richmond Road Bradford Yorks BD7 1DP	Mr I. Ashton Mr A. Diediw
University of Birmingham The Computer Centre P.O. Box 363 Birmingham B15 2TT	Mr R. Vivian
University of Bristol Computer Centre University Walk Bristol BS8 1TW	Mr C. Morris Ms S. Roberts Mr A. H. Williams

Brunel University Computer Laboratory Uxbridge Middlesex UB8 3PH	Mr N. Bonney
Cambridge University Computer Laboratory Corn Exchange Street Cambridge CB2 3QG	Mr C. J. Cheney Mr M. J. T. Guy
The City University Computer Unit St. John Street London EC1V 4PB	Mr A. D. Holt
Computer Aided Design Centre Madingley Road Cambridge CB3 ØHB	Mr A. Jones
CAP (London) Limited 233 High Holborn London WC1V 7DJ	Mr B. Wood
Science Research Council Daresbury Laboratory Daresbury Warrington Cheshire WA4 4AD	Dr P. S. Kummer
Data Communication Protocols Unit c/o Computer Laboratory Corn Exchange Street Cambridge CB2 3QG	Dr P. F. Linington
Data Communication Protocols Unit National Physical Laboratory Queens Road Teddington Middlesex TW11 ØLW	Mr K. A. Bartlett
Digital Equipment Corporation 1213 Petit-Lancy Geneva Switzerland	Mr B. Glass

University of Dundee Computing Centre Park Place Dundee DD1 4HN	Mr C. Bye
Durham University Computer Unit South Road Durham DH1 3LE	Dr P. S. Jones Mr A. A. Young
E.R.C.C. The King's Buildings Mayfield Road Edinburgh EH9 3JZ	Mr R. Chisholm Dr W. D. Hay
Essex University Computing Centre Wivenhoe Park Colchester Essex CO4 3SQ	Dr R. G. Blake Dr A. M. Stanier
University of Exeter Computer Unit Maths & Geology Building North Park Road Exeter Devon EX4 4QJ	Dr A. Buttle
GEC Computers Limited Elstree Way Borehamwood Hertfordshire WD6 1RX	Dr J. R. Horton
Glasgow University Computing Services Glasgow G11 4QQ	Mrs D. M. Dainton
Hatfield Polytechnic Computer Centre P.O. Box 1Ø9 College Lane Hatfield Herts AL1Ø 9AB	Mr R. J. Driscoll Dr M. D. Sayers

Heriot-Watt University
Computer Centre
Riccarton
Currie
Midlothian
EH14 4AS

Mr S. R. M. Clelland

Honeywell Information Systems Ltd
Data Processing Division
Windsor House
Albert Street
Slough

Mr S. Nokes

Hull University
The Computer Centre
Cottingham Road
Hull
HU6 7RX

Mr D. Spridgeon

ICL - Dataskil Limited
Reading Bridge House
Reading Bridge Approach
Reading
RG1 8PN

Mr J. Sexton

ICL - Dataskil Limited
St. Anne's House
Parsonage Green
Wilmslow
Cheshire
SK9 1HQ

Mr J. Salter

Imperial College
Computer Centre
Exhibition Road
London SW7 2BX

Mr J. Couzens

Imperial College
Dept. of Computing and Control
180 Queen's Gate
South Kensington
London SW7

Mr C. G. Pugh

Joint Network Team
c/o Rutherford Laboratory
Chilton
Didcot
Oxon
OX11 0RX

Dr B. J. Charles
Dr R. A. Rosner

University of Keele Computer Laboratory Keele Staffordshire ST5 5BG	Dr D. Wanless
University of Kent Computing Laboratory Canterbury Kent CT2 7NF	Mr S. E. Binns Mr D. J. Caul Mr I. N. Dallas Mr M. J. Garside Mr M. N. A. Lee Dr E. B. Spratt
Lancaster University Computer Services Bailrigg Lancaster LA1 4YW	Mr A. Phillips
Leeds Polytechnic Computer Unit Calverley Street Leeds LS1 3HE	Mr A. Boodson
Leeds University Computer Service Leeds LS2 9JT	Mr R. G. Powell Professor M. Wells
University of Leicester Computer Laboratory Bennett Building University Road Leicester	Mr J. E. Morris
University of Liverpool Computer Laboratory P.O. Box 147 Liverpool L69 3BX	Dr J. D. Rice
U.L.C.C. 20 Guilford Street London WC1N 1DZ	Mr A. Drabble Mr H. C. Kirkman Mr A. C. Peatfield
Loughborough University Computer Centre Ashby Road Loughborough Leicestershire	Mr J. A. Prentice

U.M.I.S.T. Control Systems Centre P.O. Box 88 Sackville Street Manchester M6Ø 1QD	Dr G. C. Barney
U.M.I.S.T. Interactive Computing Facility P.O. Box 88 Sackville Street Manchester M6Ø 1QD	Mr W. Swindells
U.M.R.C.C. Oxford Road Manchester M13 9PL	Mr V. G. Aswani Mr D. Johnson
National Physical Laboratory Division of Numerical Analysis and Computer Science Teddington Middlesex TW11 ØLW	Dr D. Rayner
Network Unit	Mr M. B. Williams
University of Newcastle Computing Laboratory Claremont Tower Claremont Road Newcastle Upon Tyne NE1 7RU	Mr J. P. Aspden Mr F. T. Greenwood
North London Polytechnic Computing Services Holloway Road London N7 8DB	Mr A. Austen Mr G. Jacob
Nottingham University Cripps Computing Centre University Park Nottingham NG7 2RD	Mr P. S. Harrison Mr A. Jack

Institute of Oceanographic Sciences Bidston Observatory Bidston Birkenhead Merseyside L43 7RA	Mr D. J. Hughes
Oxford University Computing Service 13 Banbury Road Oxford	Mr G. W. Litchfield
Post Office Telecommunications Lutyens House 1-6 Finsbury Circus London EC2M 7LY	Mr P. Morrison
Post Office Telecommunications River Plate House Finsbury Circus London EC2M 7LY	Mr M. Sands
Prime Computers Inc. Research and Development Boston Massachusetts U.S.A.	Mr Gurteen
Prime Computers Research Limited The Merton Centre St. Peters Street Bedford	Mr T. L. Wallace
Queen Mary College Computer Centre Mile End Road London E1 4NS	Mr J. P. Brandon
Queen's University Computer Unit Belfast BT7 1NN	Mr D. O'Sullivan
University of Reading Computer Centre Whiteknights Reading Berkshire RG6 2AX	Mr N. P. C. Lynch

Rothamsted Experimental Station Computer Unit Harpenden Hertfordshire AL5 2JQ	Mr I. M. Richmond
Science Research Council Rutherford Laboratory Computer Unit Chilton Didcot Oxon OX11 ØQX	Mr P. E. Bryant Mr J. Burren Dr C. Cooper Mr P. Girard
University of St. Andrews Computing Laboratory North Haugh St. Andrews Fife KY16 9SX	Mr M. Piotrowicz
University of Salford Computing Laboratory The Crescent Salford M5 4WT	Mr H. D. Ellison Dr J. Larmouth
Sheffield University Computing Services Sheffield S1Ø 2TN	Mr P. Mason
Southampton University Computing Service Southampton SO9 5NH	Mr M. W. Sage
South Bank Polytechnic Computer Centre Borough Road London SE1 ØAA	Mr S. Wanless Mr P. Watts
South West Universities Computer Network University of Exeter Maths & Geology Building North Park Road Exeter Devon EX4 4QJ	Dr H. E. Davies

S.W.U.R.C.C. University of Bath Claverton Down Bath BA2 7AY	Mr R. G. Bradshaw Mr A. K. Dand Mr J. S. Thomas
University of Stirling Computer Unit Stirling FK9 4LA	Mr A. J. Scolley
Strathclyde University Computer Laboratory 204 George Street Glasgow G1 1XW	Mr J. B. Jamieson
Teesside Polytechnic Computer Centre Borough Road Middlesborough TS1 3BA	Dr P. Barkataki
Trinity College Department of Computer Science Dublin 2 Eire	Mr A. M. Patel
The New University of Ulster Computer Centre Coleraine Co. Londonderry	Mr R. M. Gillman
University College, London Computer Unit 17-19 Gordon Street London WC1H 0AH	Mr C. J. Kennington
University College, London Statistics & Computer Science Department Gower Street London WC1E 6BT	Mr P. L. Higginson Mr A. Hinchley Mr S. R. Wilbur
University College, Swansea Computer Centre Singleton Park Swansea SA2 8PP	Mr P. R. Brady

U.W.I.S.T.
Computer Department
King Edward VII Avenue
Cardiff
CF1 3NU

Dr D. Jennings

University of Warwick
Computer Unit
Coventry
CV4 7AL

Mr T. B. G. Clark

Westfield College
Kidderpore Avenue
London NW3 7ST

Mr M. G. Johnson

University of York
Department of Computer Science
Heslington
York
YO1 5DD

Mr J. D. Service